



US009128995B1

(12) **United States Patent**  
**Fletcher et al.**

(10) **Patent No.:** **US 9,128,995 B1**  
(45) **Date of Patent:** **Sep. 8, 2015**

(54) **DEFINING A GRAPHICAL VISUALIZATION  
ALONG A TIME-BASED GRAPH LANE  
USING KEY PERFORMANCE INDICATORS  
DERIVED FROM MACHINE DATA**

7,800,613 B2 \* 9/2010 Hanrahan et al. .... 345/440  
7,848,260 B2 \* 12/2010 Cowan ..... 370/255  
8,095,417 B2 1/2012 Handy et al.  
8,266,148 B2 9/2012 Guha et al.  
8,364,460 B2 1/2013 Ostermeyer et al.  
8,412,696 B2 4/2013 Zhang et al.  
8,538,787 B2 9/2013 Braun et al.

(71) Applicant: **Splunk Inc.**, San Francisco, CA (US)

(Continued)

(72) Inventors: **Tristan Antonio Fletcher**, Pacifica, CA  
(US); **Alok Anant Bhide**, Mountain  
View, CA (US)

#### OTHER PUBLICATIONS

(73) Assignee: **Splunk, Inc.**, San Francisco, CA (US)

Jack Coates, Cognitive Splunking, Sep. 17, 2012; Splunk-blogs,  
Blogs-Security, 1-3.\*

(\*) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 0 days.

(Continued)

(21) Appl. No.: **14/528,978**

*Primary Examiner* — Pavan Mamillapalli

(22) Filed: **Oct. 30, 2014**

(74) *Attorney, Agent, or Firm* — Law Office of Thomas L.  
Treffert

#### Related U.S. Application Data

(60) Provisional application No. 62/062,104, filed on Oct.  
9, 2014.

(51) **Int. Cl.**  
**G06F 17/30** (2006.01)  
**G06F 7/00** (2006.01)

(52) **U.S. Cl.**  
CPC .... **G06F 17/30554** (2013.01); **G06F 17/30958**  
(2013.01)

(58) **Field of Classification Search**  
CPC . G06F 11/3409; G06F 11/30; G06F 11/3466;  
G06F 3/04842; G06F 8/34; G06F 17/30554;  
G06F 17/30958  
USPC ..... 707/722  
See application file for complete search history.

(56) **References Cited**


#### U.S. PATENT DOCUMENTS

7,299,358 B2 11/2007 Chateau et al.  
7,711,670 B2 5/2010 Roediger  
7,778,952 B2 \* 8/2010 Vespe et al. .... 706/56

(57) **ABSTRACT**

A service monitoring system receives a selection of key performance indicators (KPIs) that each indicate a different aspect of how a service provided by one or more entities is performing. Each entity of the one or more entities produces machine data or wherein each entity has its operation reflected in machine data not produced by the entity. Each KPI is defined by a different search query that derives one or more values from the machine data pertaining to the one or more entities providing the service, where each of the one or more values is associated with a point in time and representing the aspect of how the service is performing at the associated point in time. For each of the selected KPIs, the service monitoring system derives the one or more values and causes display of a graphical visualization of the derived one or more values for the KPI along a time-based graph lane. The graph lanes for the selected KPIs are parallel to each other and the graphical visualizations in the graph lanes are all calibrated to a same time scale.

**30 Claims, 91 Drawing Sheets**



(56)

**References Cited****U.S. PATENT DOCUMENTS**

8,543,527 B2 9/2013 Bates et al.  
 8,589,403 B2 11/2013 Marquardt et al.  
 8,682,925 B1 3/2014 Marquardt et al.  
 8,712,953 B2 4/2014 Beringer et al.  
 8,738,414 B1\* 5/2014 Nagar et al. .... 705/7.25  
 8,806,361 B1 8/2014 Noel et al.  
 8,825,752 B1 9/2014 Madhavan  
 8,948,369 B2 2/2015 Shaffer et al.  
 2001/0049682 A1 12/2001 Vincent et al.  
 2003/0174173 A1 9/2003 Nishiyama et al.  
 2003/0182310 A1 9/2003 Charnock et al.  
 2004/0024770 A1\* 2/2004 Cardno ..... 707/100  
 2004/0030668 A1 2/2004 Pawlowski et al.  
 2005/0060048 A1\* 3/2005 Pierre et al. .... 700/28  
 2005/0181835 A1 8/2005 Lau et al.  
 2006/0156250 A1 7/2006 Chaudhri et al.  
 2007/0005388 A1 1/2007 Busch et al.  
 2007/0150480 A1 6/2007 Hwang et al.  
 2007/0192150 A1 8/2007 Belkin et al.  
 2007/0208601 A1\* 9/2007 Pulianda ..... 705/7  
 2008/0081632 A1 4/2008 Malik  
 2008/0097807 A1 4/2008 Chang et al.  
 2008/0120129 A1 5/2008 Seubert et al.  
 2008/0126417 A1 5/2008 Mazurik  
 2008/0140514 A1 6/2008 Stenger  
 2008/0177595 A1 7/2008 Wu et al.  
 2008/0201397 A1\* 8/2008 Peng et al. .... 708/308  
 2008/0256516 A1 10/2008 Chaar et al.  
 2009/0112932 A1 4/2009 Skierkowski et al.  
 2009/0265637 A1 10/2009 Lee et al.  
 2009/0313503 A1 12/2009 Atluri et al.  
 2010/0023362 A1 1/2010 Nguyen et al.  
 2010/0031234 A1 2/2010 Chaar et al.  
 2010/0042680 A1 2/2010 Czyzewicz et al.  
 2010/0324927 A1 12/2010 Tinsley  
 2010/0324962 A1 12/2010 Nesler et al.  
 2010/0332466 A1 12/2010 White et al.  
 2011/0178977 A1 7/2011 Drees  
 2011/0261055 A1 10/2011 Wong et al.  
 2011/0313817 A1 12/2011 Wang  
 2012/0005593 A1 1/2012 Redpath

2012/0158521 A1 6/2012 McCullen  
 2012/0162265 A1 6/2012 Heinrich et al.  
 2012/0259583 A1 10/2012 Noboa et al.  
 2013/0142322 A1 6/2013 Grasso et al.  
 2013/0182700 A1 7/2013 Figura et al.  
 2013/0185693 A1 7/2013 Chaar et al.  
 2013/0318236 A1 11/2013 Coates et al.  
 2013/0318589 A1 11/2013 Ford et al.  
 2013/0318603 A1 11/2013 Merza  
 2013/0325147 A1 12/2013 Karnouskos  
 2013/0326620 A1 12/2013 Merza et al.  
 2013/0332472 A1\* 12/2013 Vogel et al. .... 707/754  
 2014/0040306 A1 2/2014 Gluzman Peregrine et al.  
 2014/0067836 A1\* 3/2014 Holmes et al. .... 707/756  
 2014/0072115 A1 3/2014 Makagon et al.  
 2014/0129298 A1 5/2014 Hulen et al.  
 2014/0146648 A1 5/2014 Alber et al.  
 2014/0157142 A1 6/2014 Heinrich et al.  
 2014/0160238 A1 6/2014 Yim et al.  
 2014/0177819 A1 6/2014 Vymenets et al.  
 2014/0236889 A1 8/2014 Vasan et al.  
 2014/0236890 A1 8/2014 Vasan et al.  
 2014/0324448 A1\* 10/2014 Lacy et al. .... 705/2  
 2014/0337871 A1 11/2014 Garcia De Blas et al.  
 2014/0376710 A1 12/2014 Shaffer et al.  
 2015/0112700 A1\* 4/2015 Sublett et al. .... 705/2

**OTHER PUBLICATIONS**

Bitincka, Ledion, et al., "Optimizing Data Analysis with a Semi-Structured Time Series Database", Splunk Inc., 2010 pp. 1-9.  
 Carasso, David, "Exploring Splunk Search Processing Language (SPL) Primer and Cookbook", Splunk Inc., 2012 CITO Research, New York, 156 Pages.  
[http://docs.splunk.com/Documentation/PCI/2.1.1/\[000119\]](http://docs.splunk.com/Documentation/PCI/2.1.1/[000119]) User/IncidentReviewdashboard, 2 Pages (Last accessed Aug. 5, 2014).  
 "vSphere Monitoring and Performance", VMware, Inc., Update 1, vSphere 5.5, EN-001357-02, 2010-2014, pp. 1-174 <http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-551-monitoring-performance-guide.pdf>.  
 U.S. Appl. No. 14/167,316, filed Jan. 29, 2014.  
 U.S. Appl. No. 14/448,995, filed Jul. 31, 2014.  
 U.S. Appl. No. 14/326,459, filed Jul. 31, 2014.

\* cited by examiner

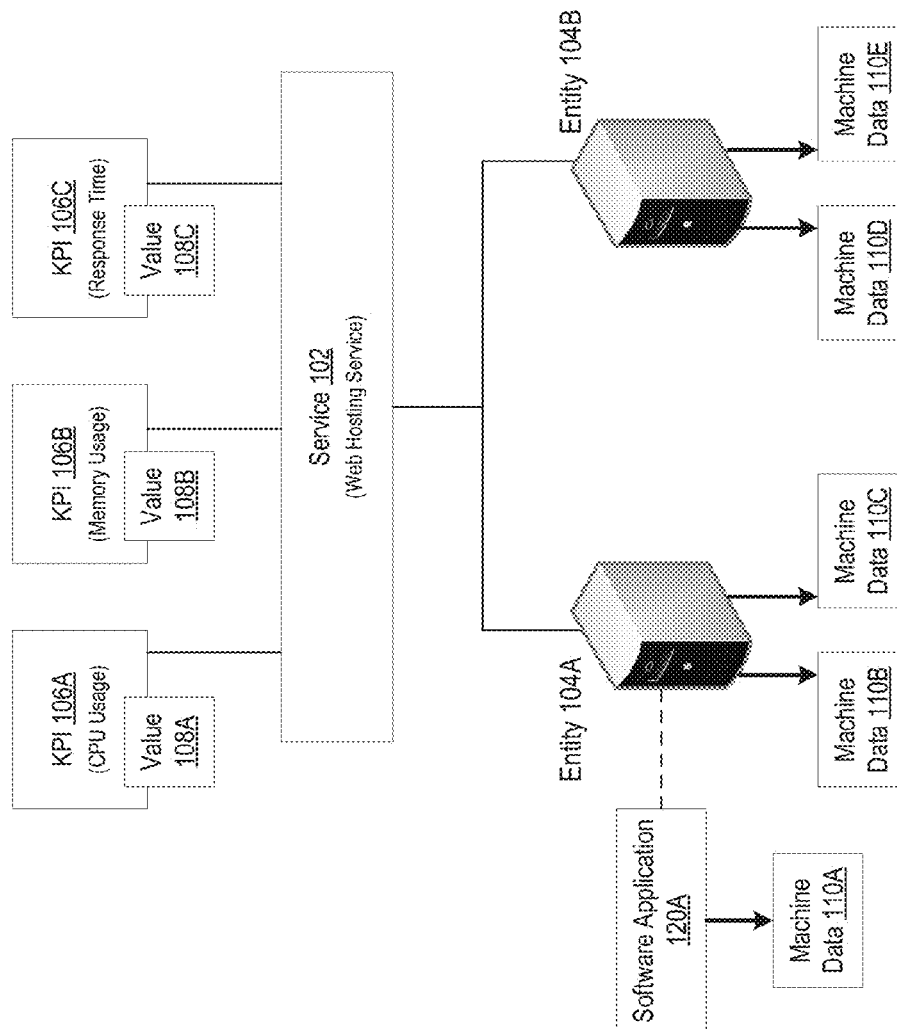


FIG. 1

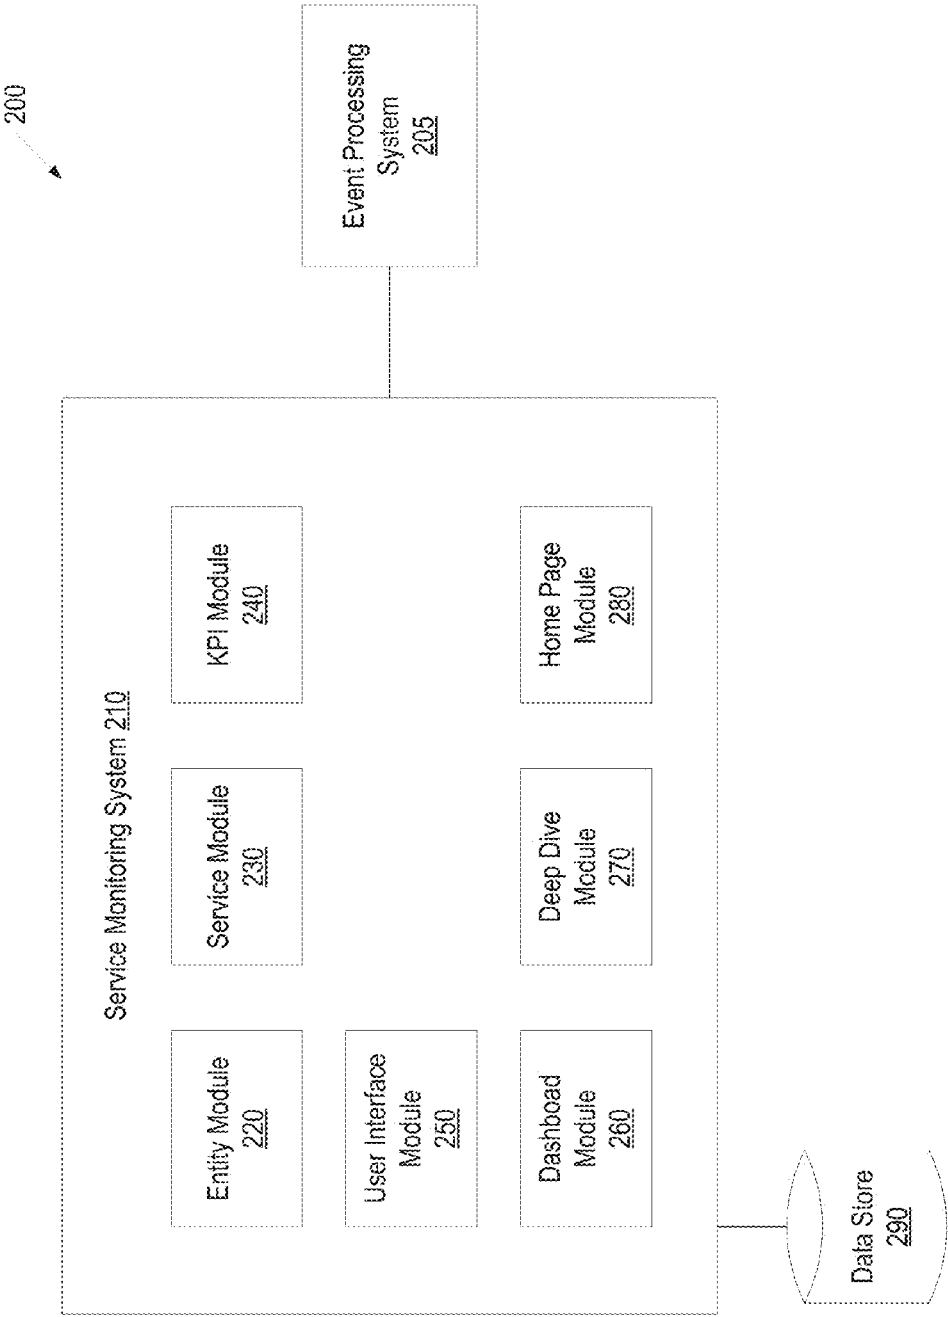
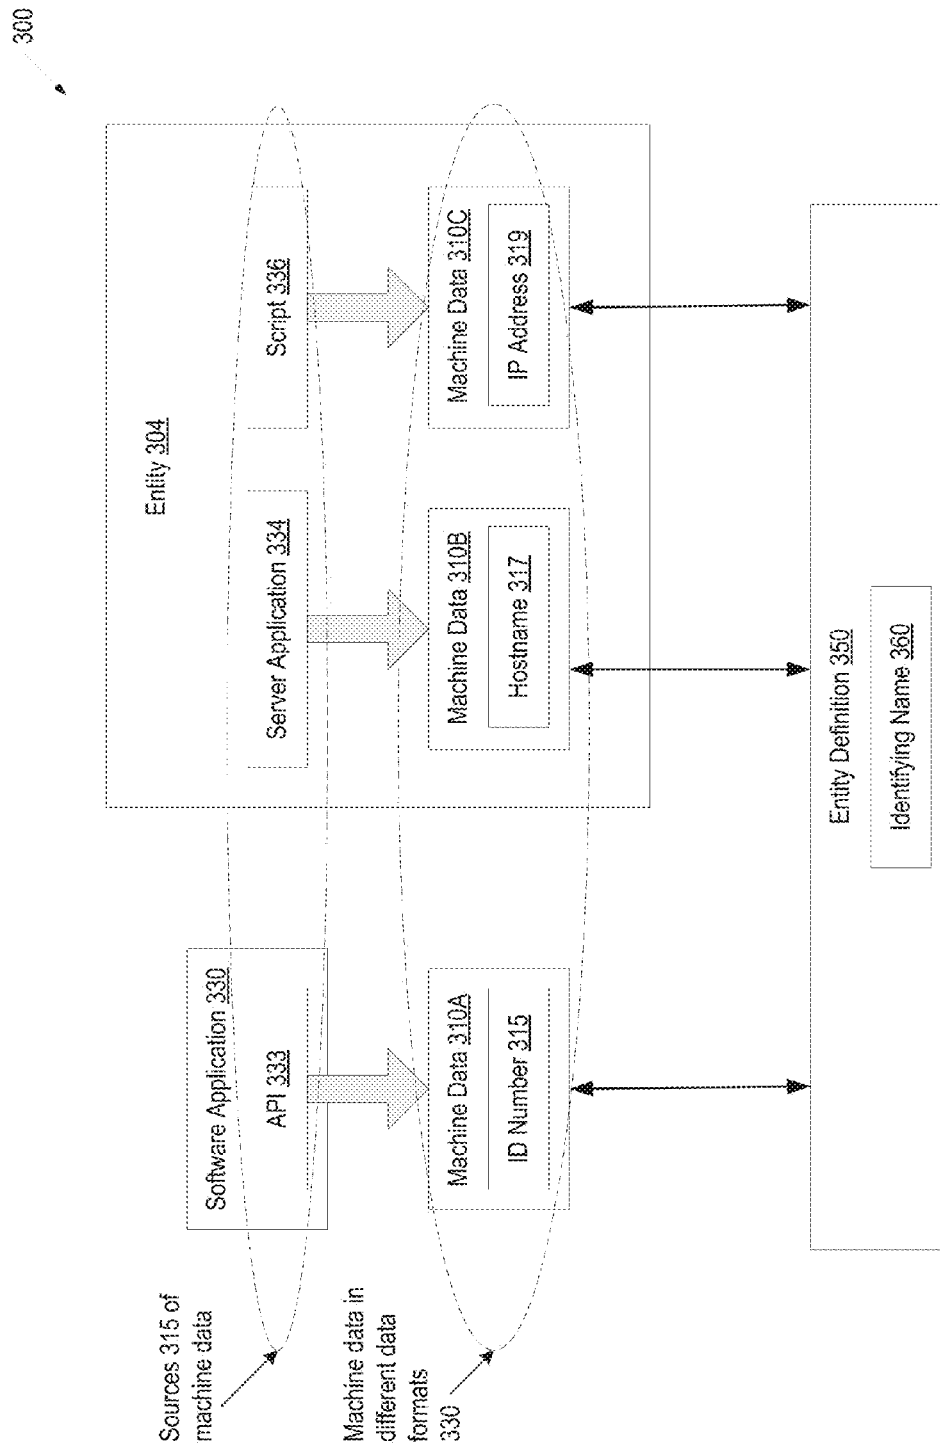


FIG. 2





366




FIG. 4

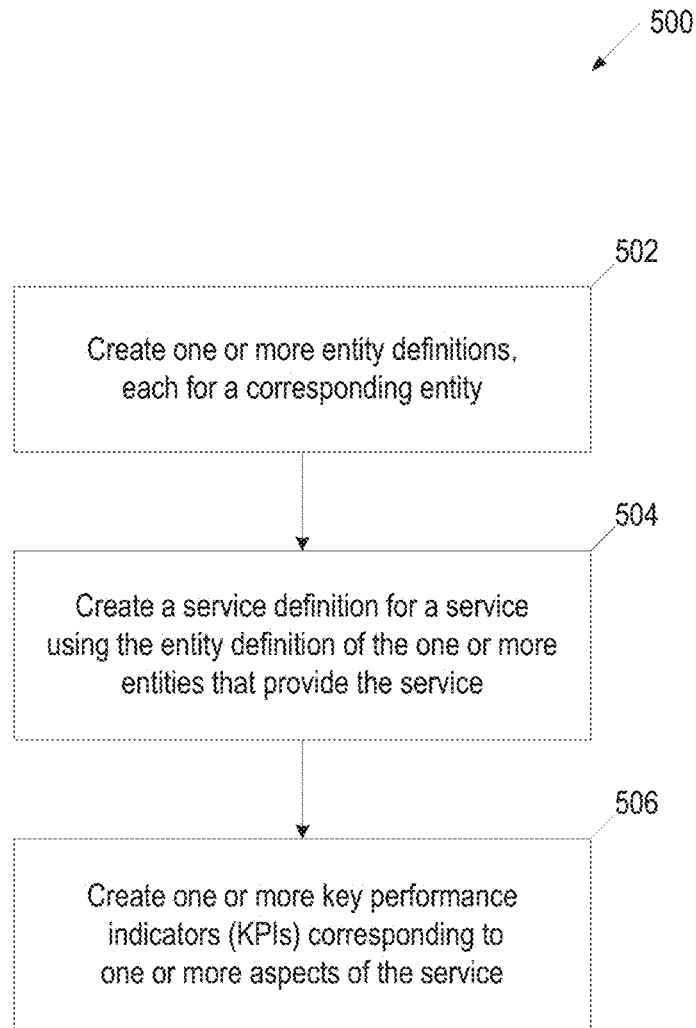


FIG. 5

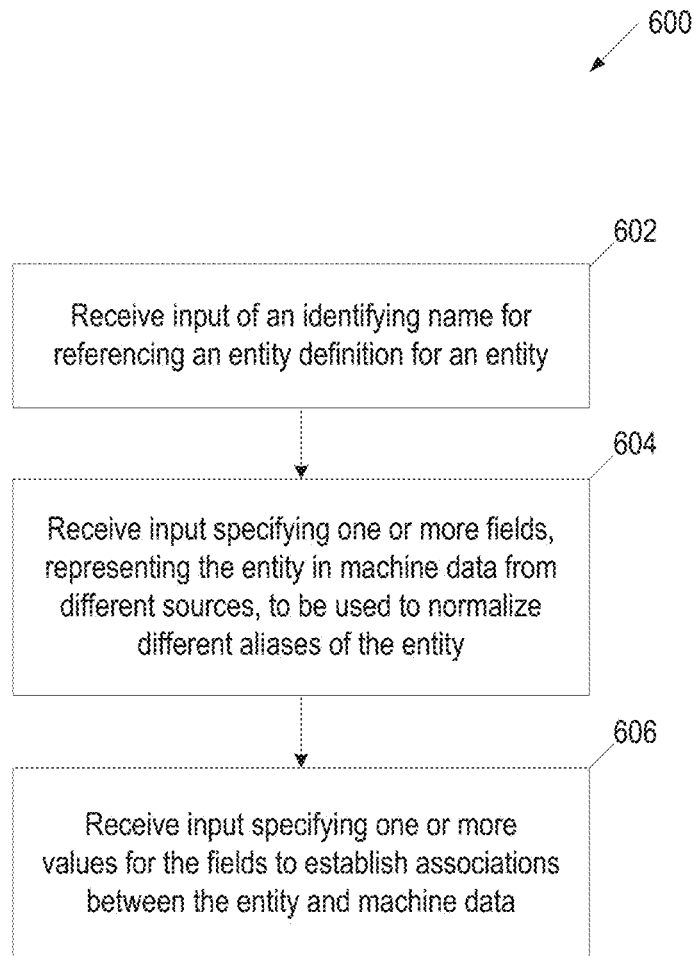




FIG. 6



800

FIG. 8



900



Entity Definition

Entity Name

Entity Type 904


Entity ID 906

Entity Description 908

Entity Address 910

Entity 912

FIG. 9A



১৯৬৫



1000




FIG. 10

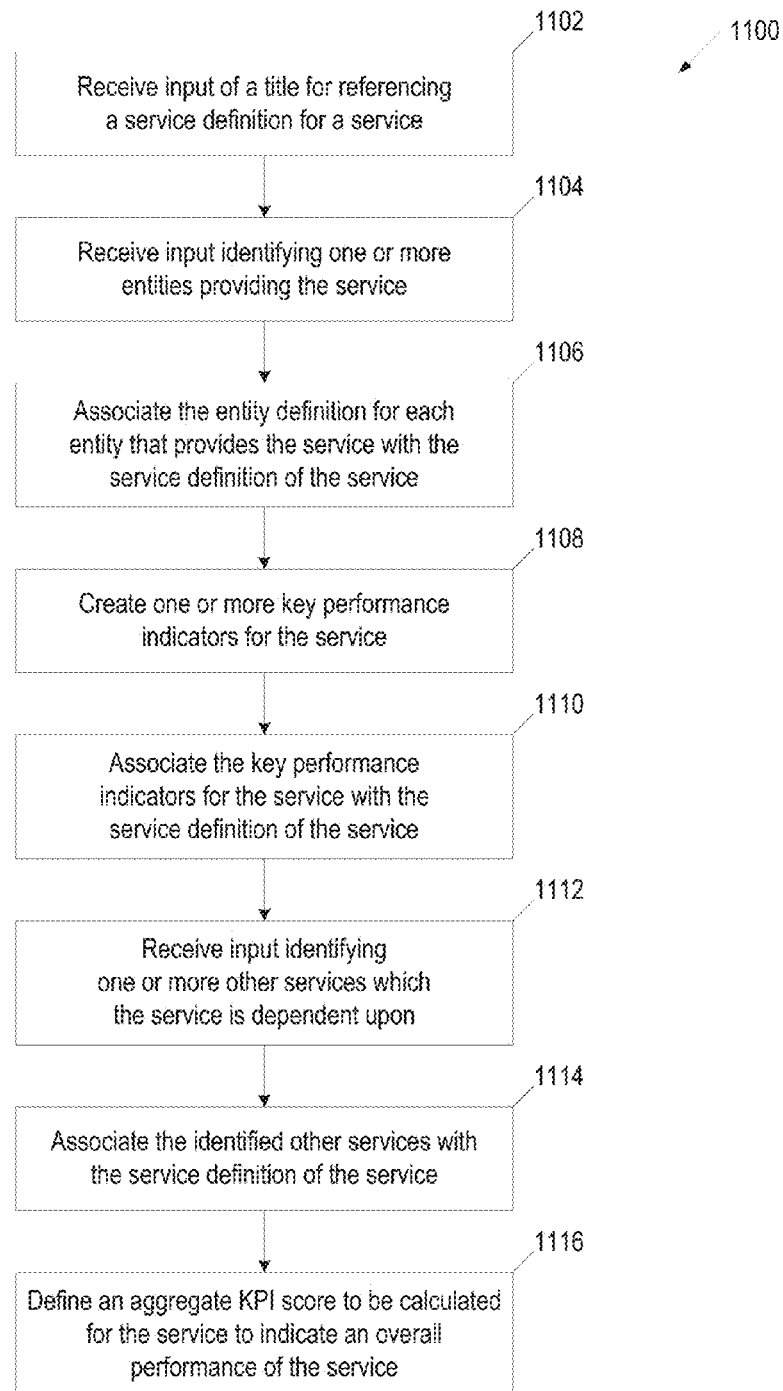


FIG. 11




FIG. 12

1300

1302

1304

1306

Create New Service Definition

Title

Web Hosting

Description

optional

Cancel

CONTINUE

FIG. 13




FIG. 14




FIG. 15

1600

Create New Entity Definition

X

Create from

New 1603

Import CSV 1605

Saved Search

Name

Description

Type

Search Fields

Search Values

1601

Close

Save changes

FIG. 16

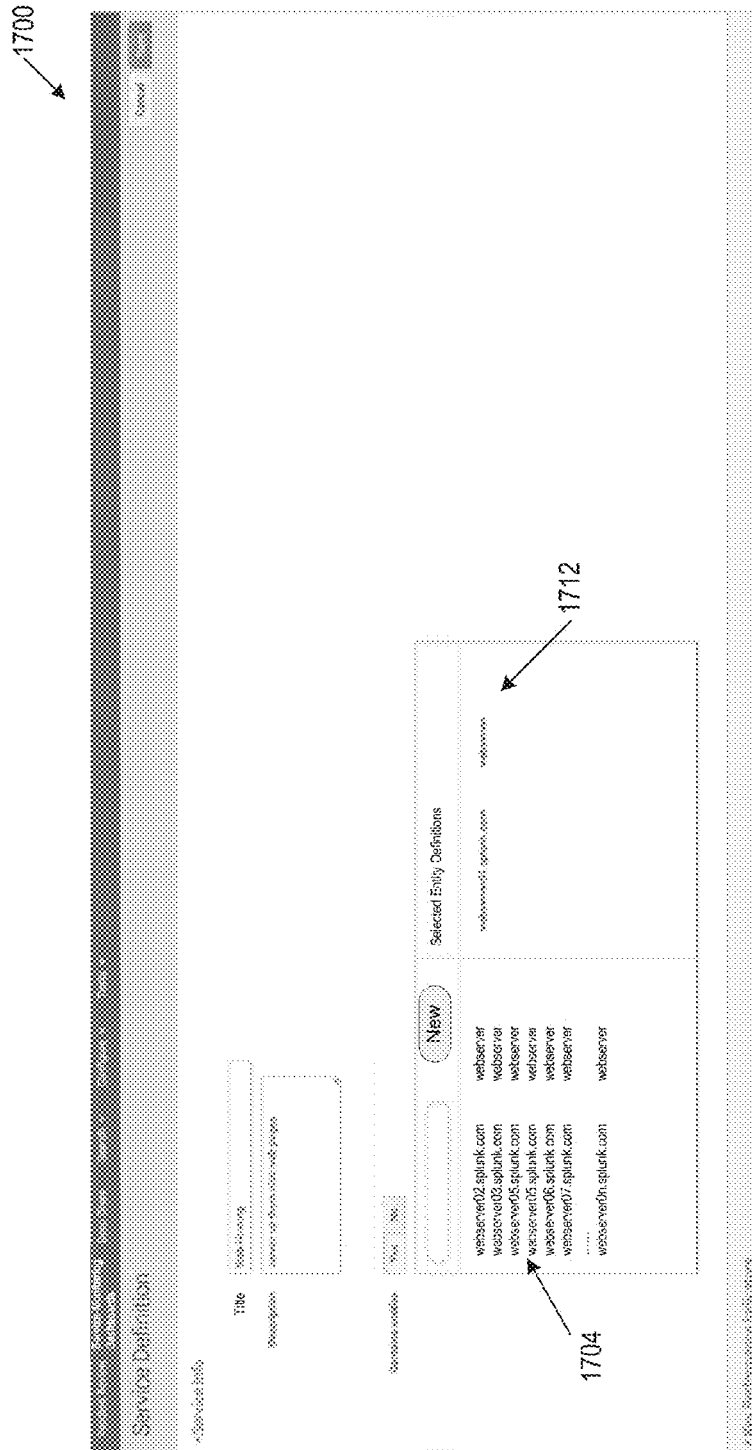


FIG. 17






FIG. 18

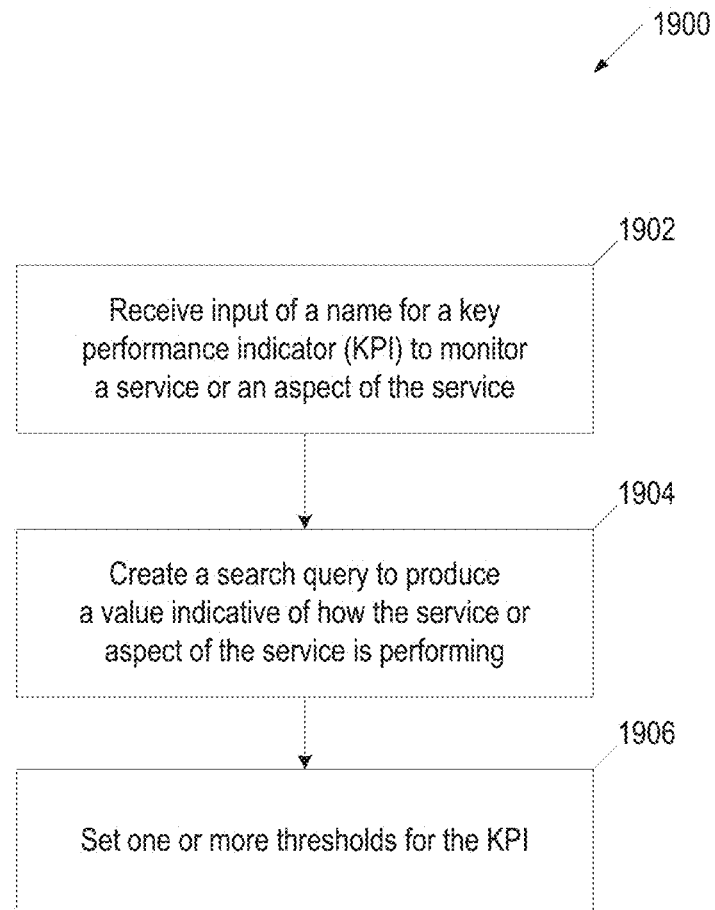


FIG. 19

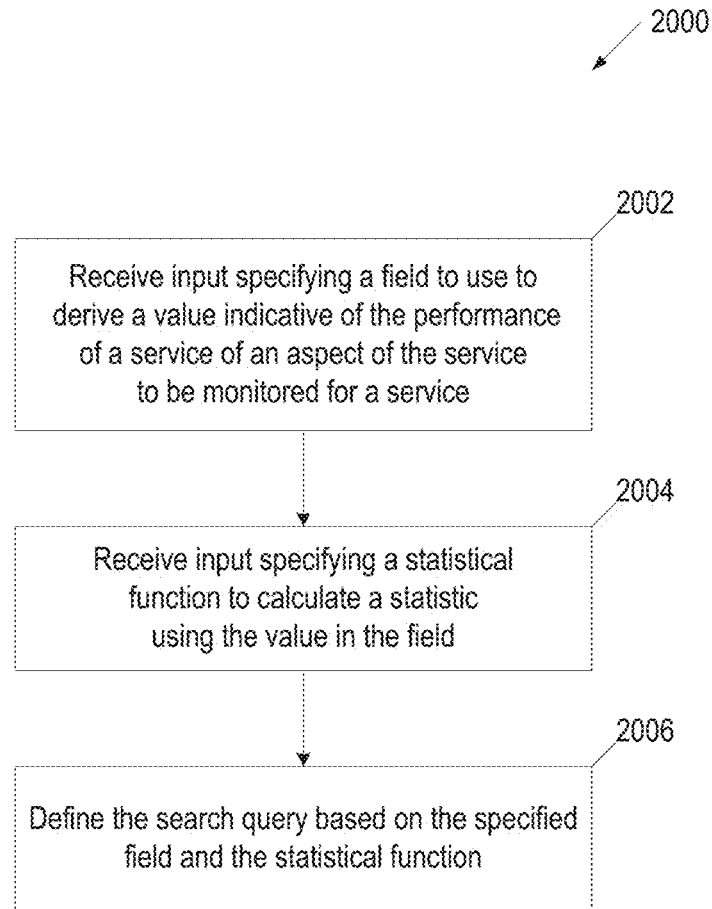


FIG. 20




FIG. 21




FIG. 22




FIG. 23

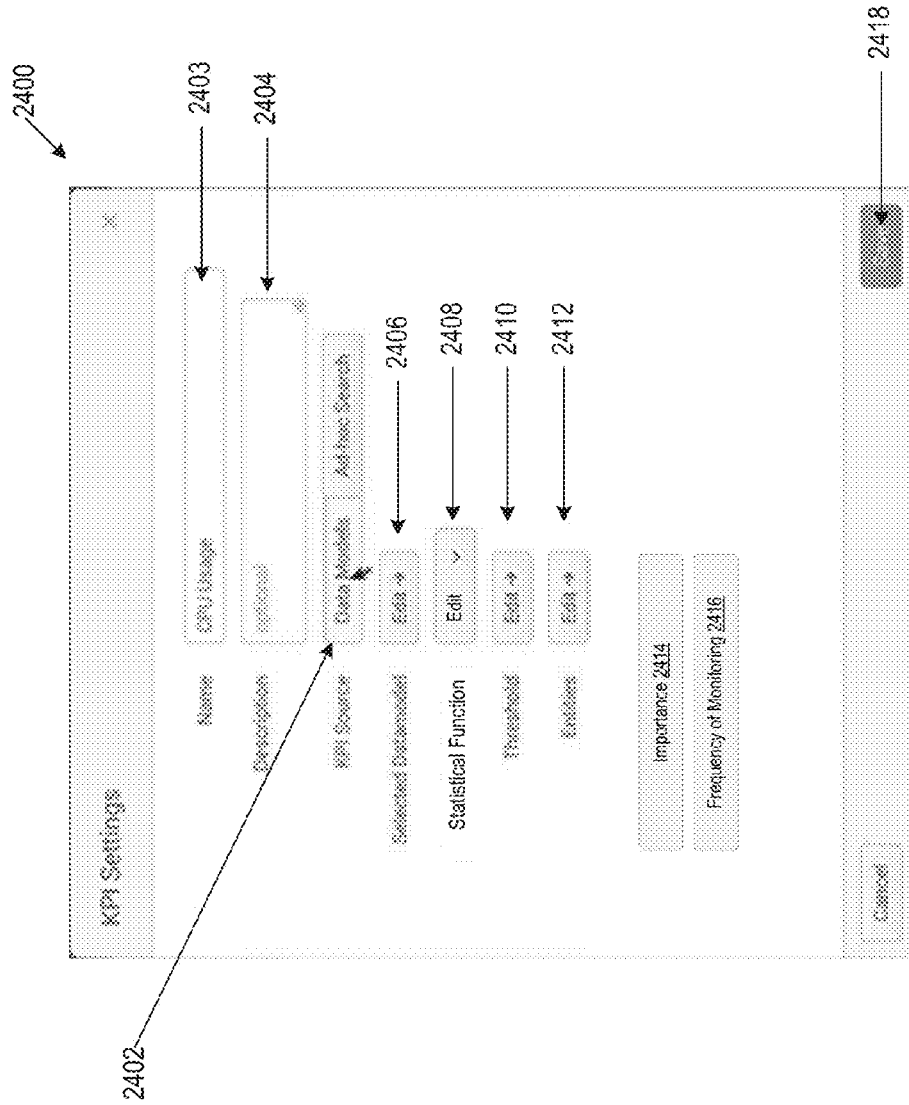


FIG. 24

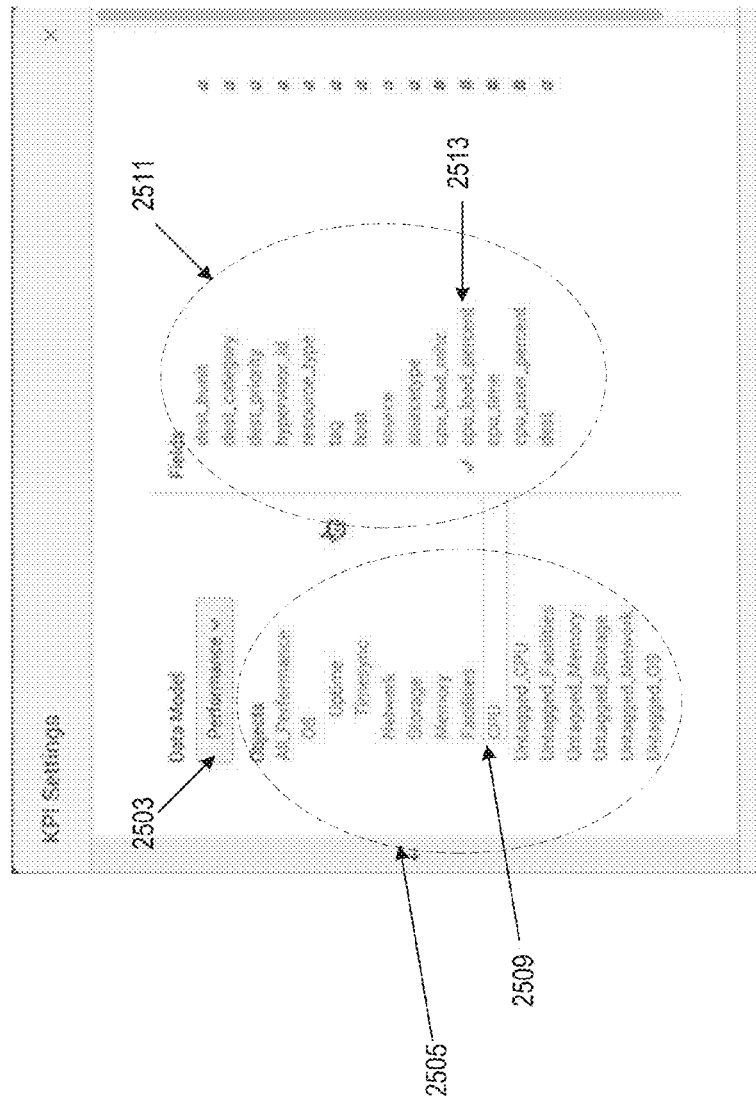


Fig. 25



2600

KPI Settings

Name: CPU Usage

Description: optional

KPI Source: Custom Models Address Search

Selected Data Model: E28

Statistical Function: Average

Threshold: E28

Statistics: E28

Importance:

Frequency of Monitoring:

Cancel

Save

2601

FIG. 26

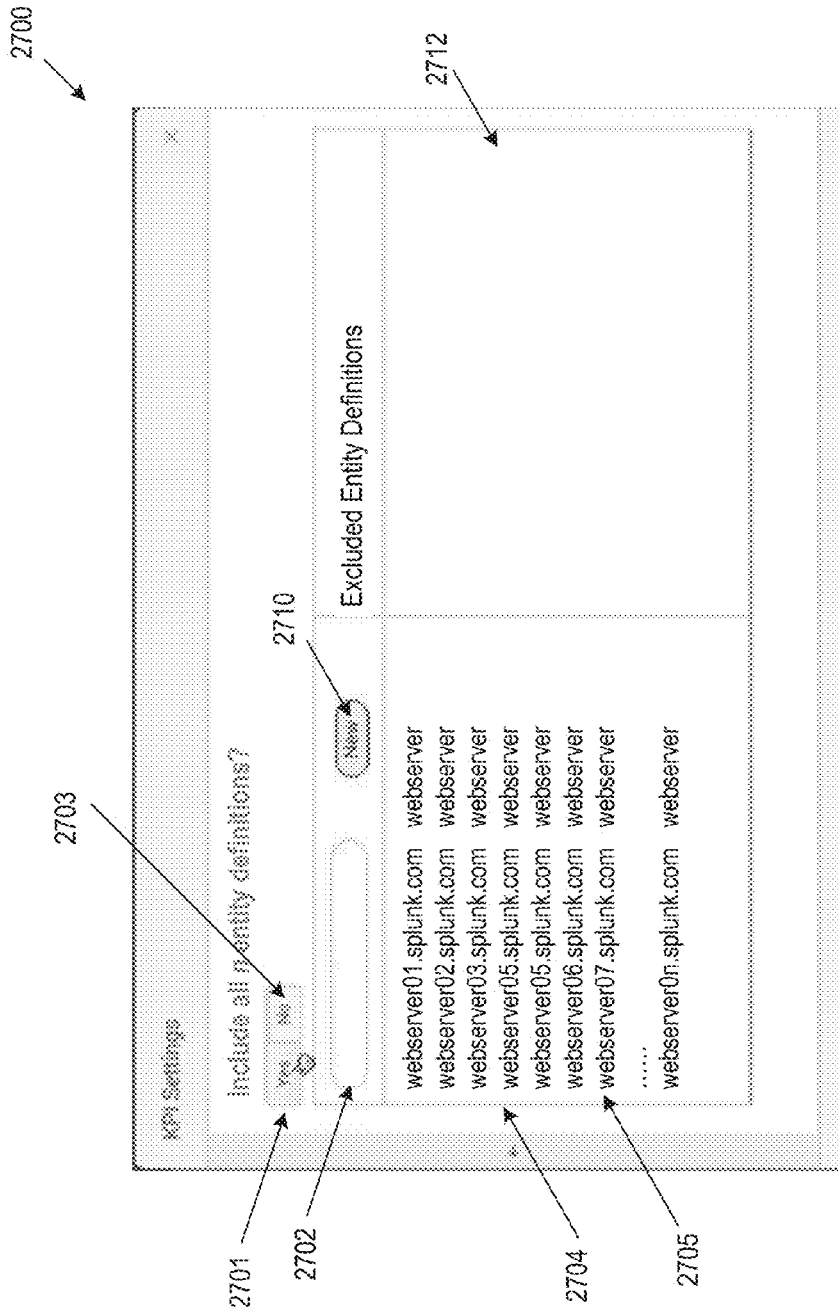


FIG. 27

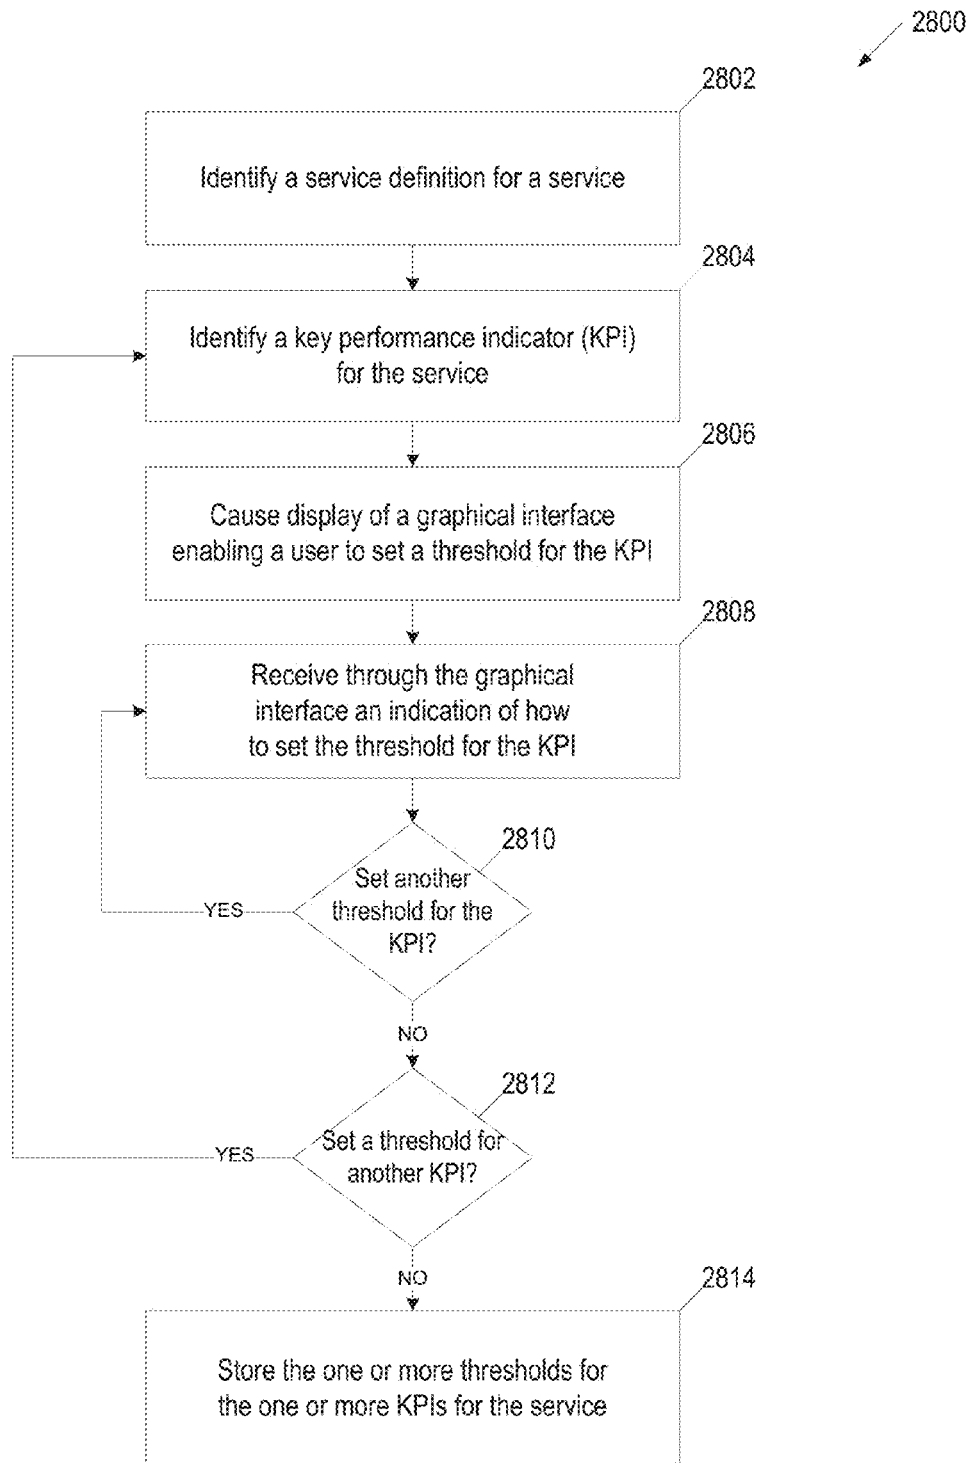


FIG. 28

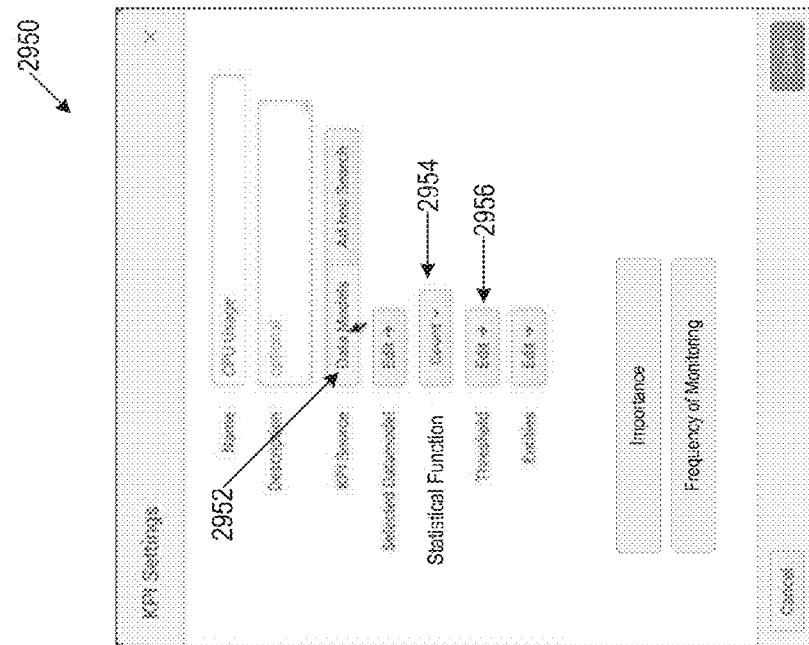


FIG. 29B

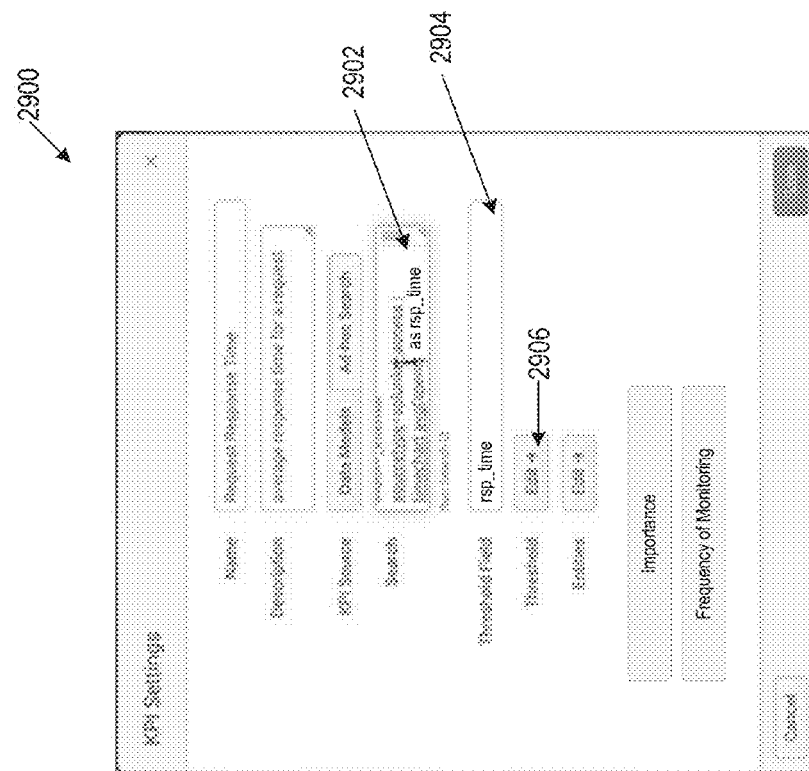


FIG. 29A

3000




FIG. 30

3100




FIG. 31A

3150




FIG. 31B

3160



FIG. 31C



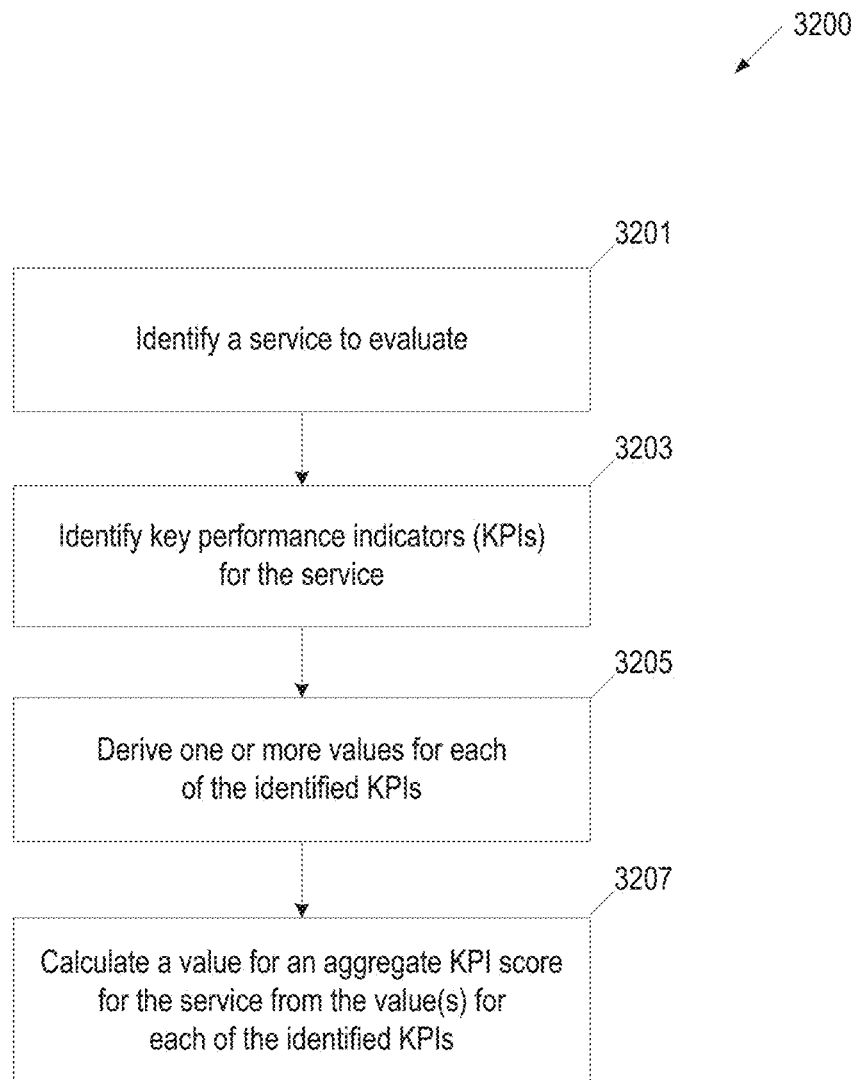


FIG. 32

3300

×

KPI Settings

Name

CPU Usage

Description

optional

KPI Source

Cpu Monitors Add New Search

Selected Data Model

Edit →

Statistical Function

Average ▾

Threshold

Edit →

Entities

Edit →

Importance 3309

Frequency of Monitoring 3311

Cancel

3313

FIG. 33A

3350




FIG. 33B

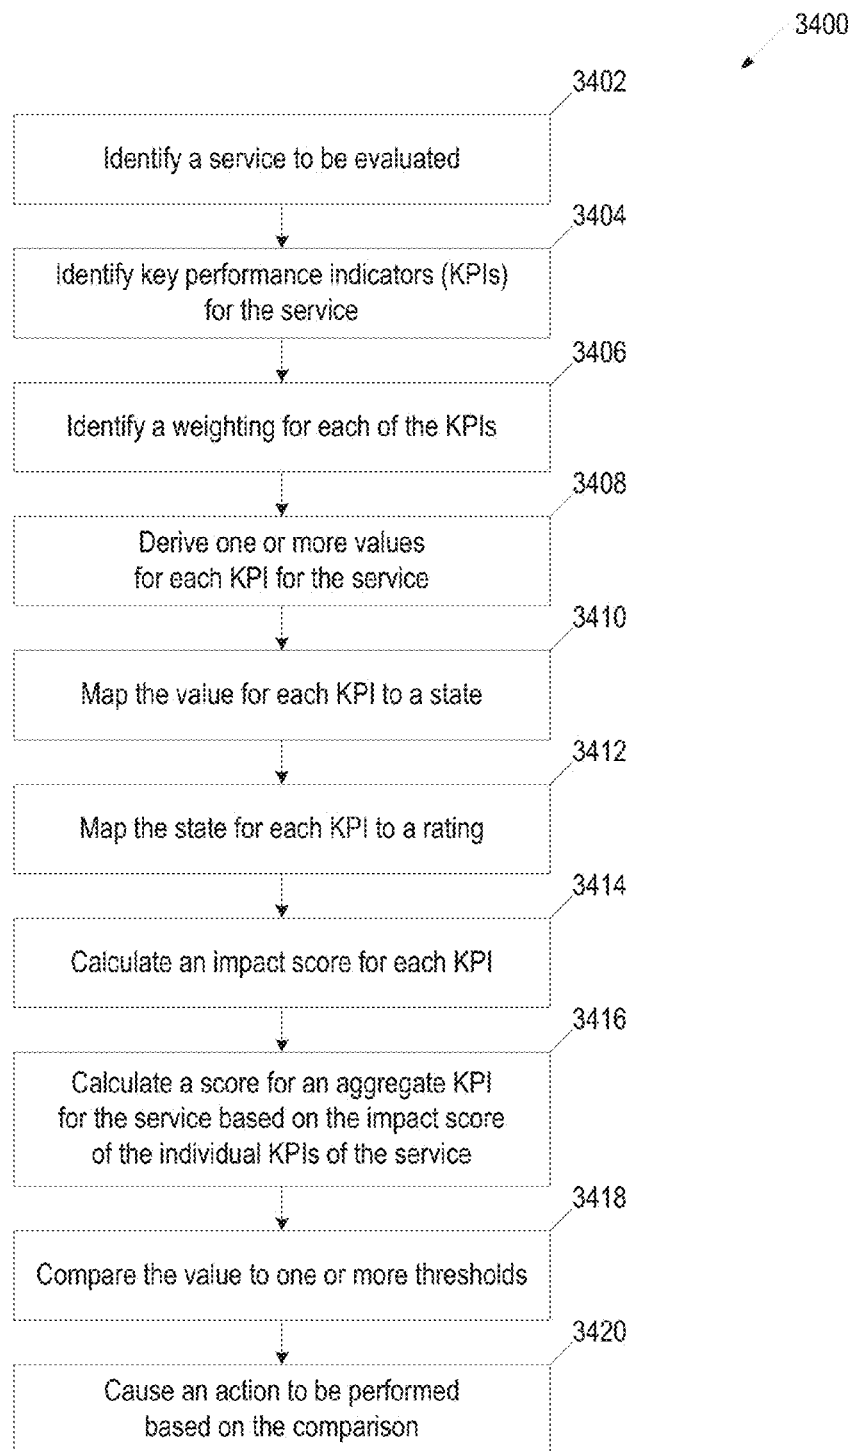


FIG. 34

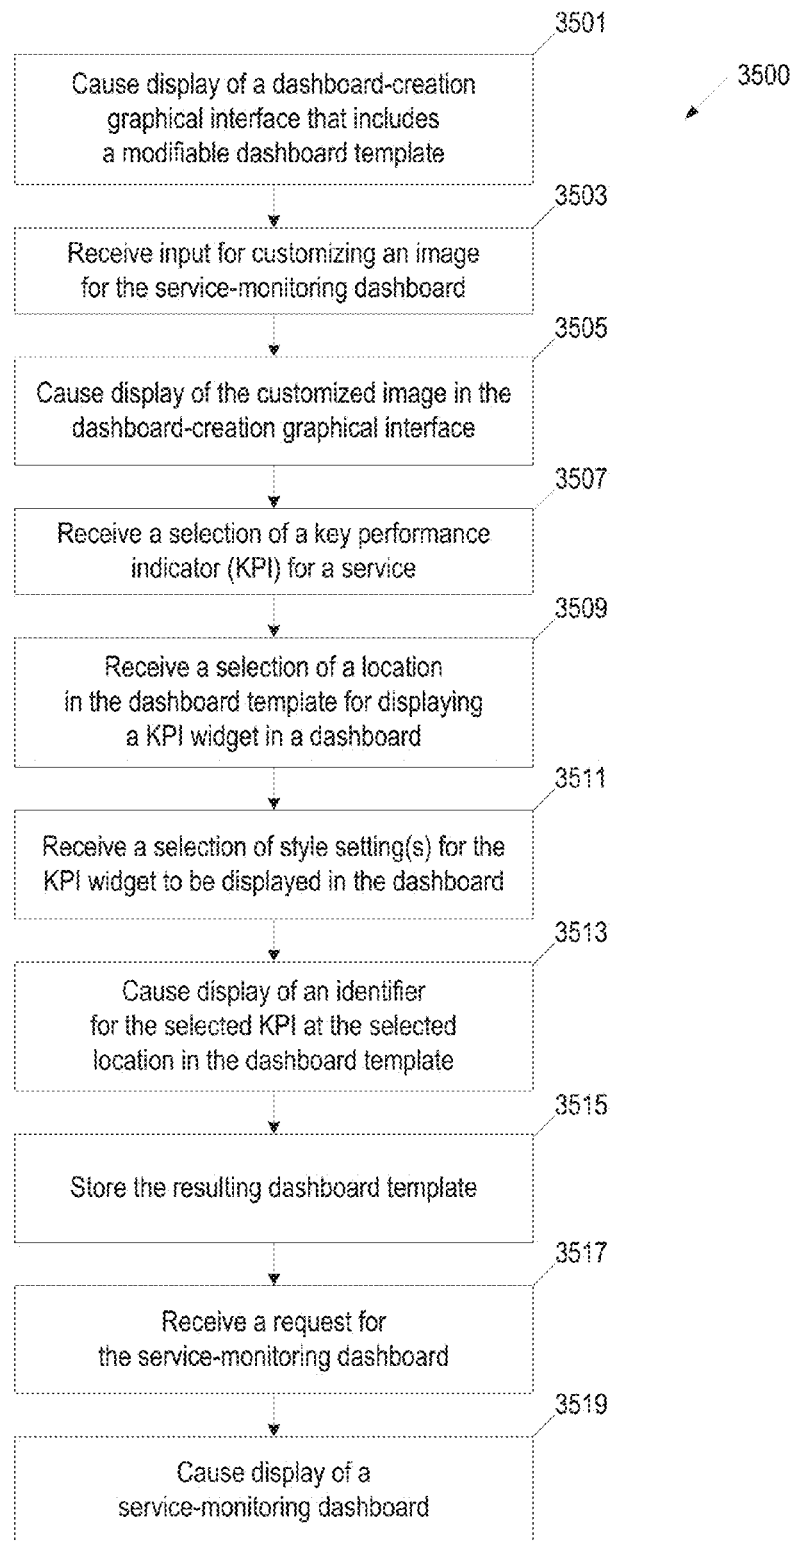


FIG. 35




FIG. 36A




FIG. 36B




FIG. 37



3800




FIG. 38




FIG. 39




FIG. 41




FIG. 40




FIG. 42

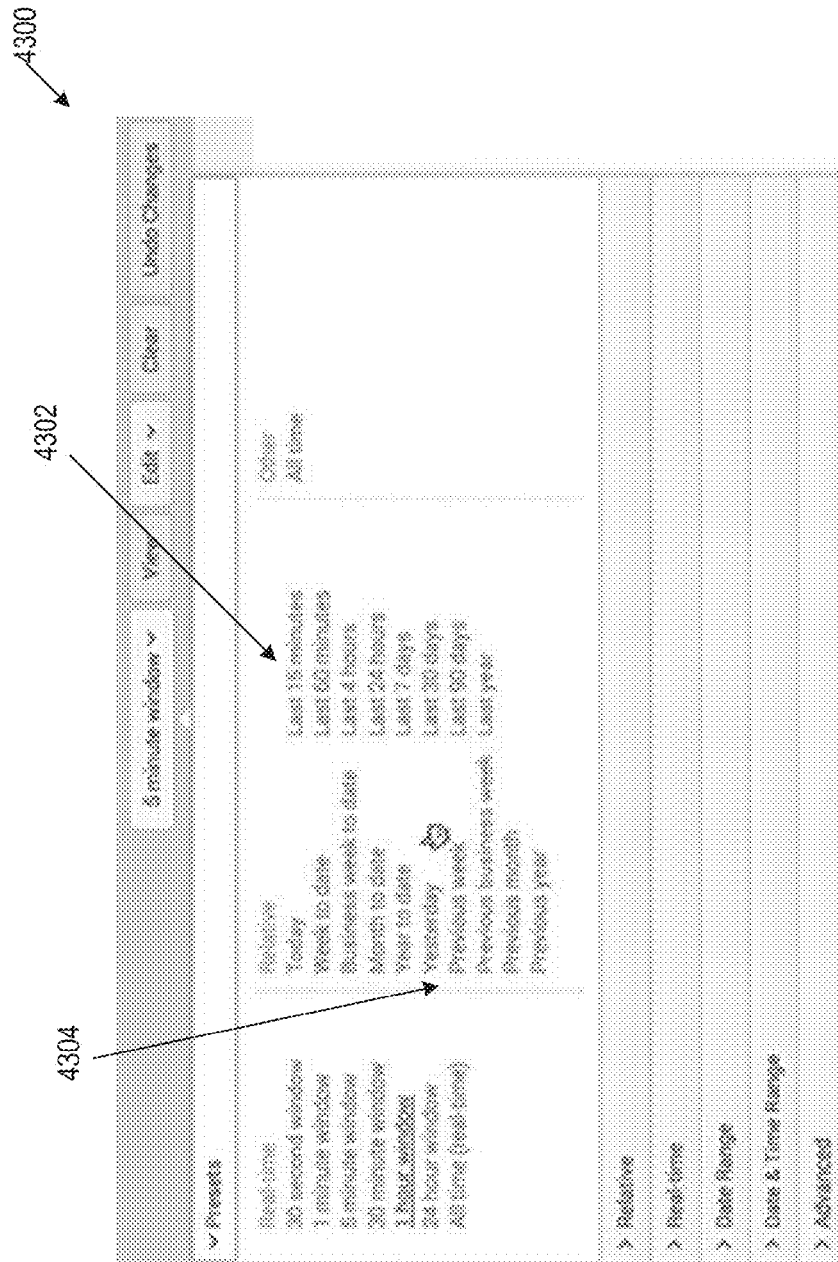




FIG. 43

FIG. 44






FIG. 46

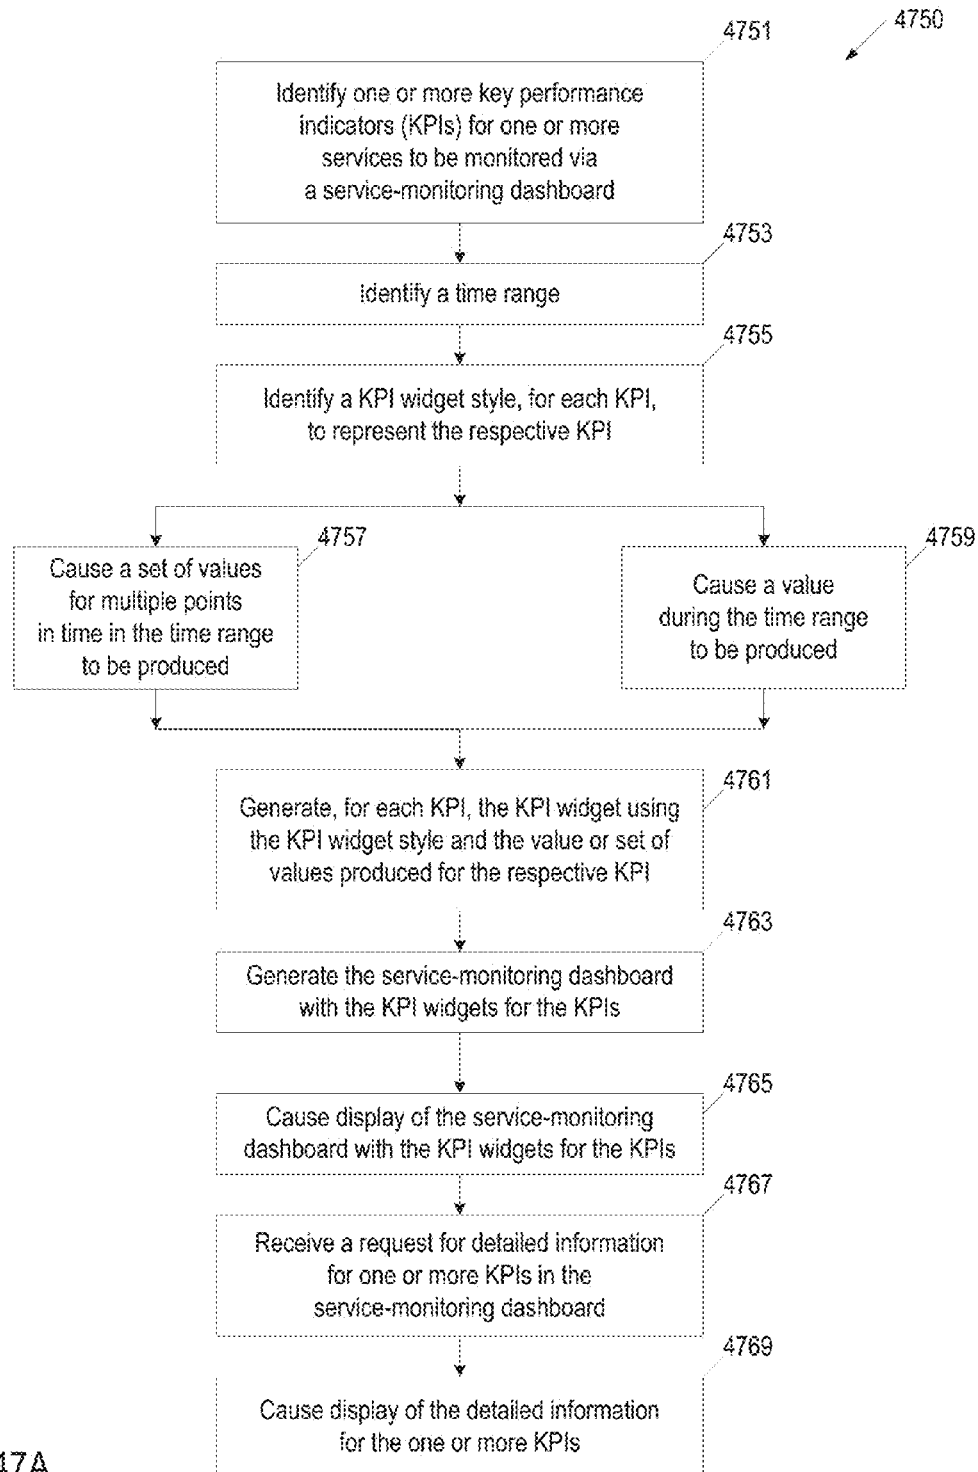


FIG. 47A

4700


4702

4708

4718

4712

FIG. 47B





4800

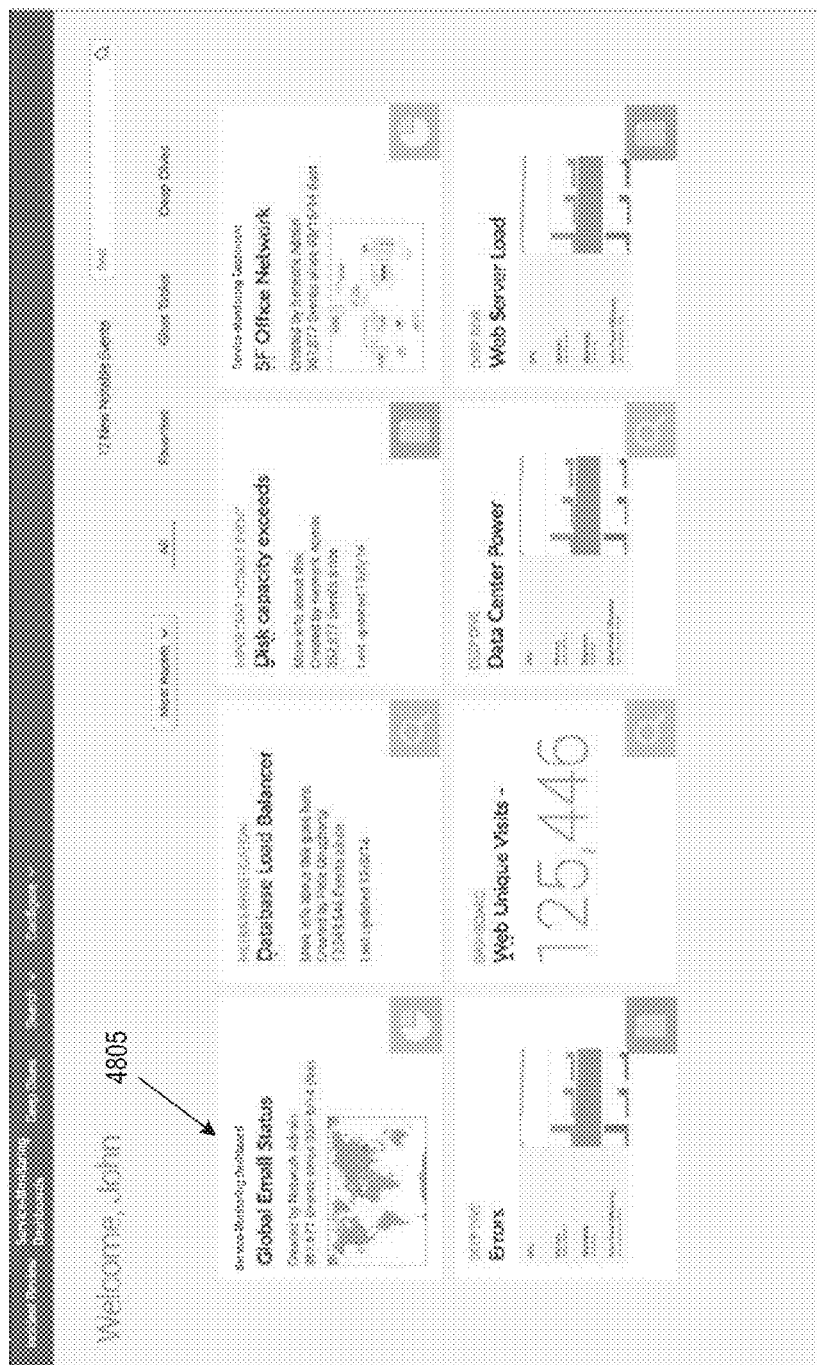


FIG. 48

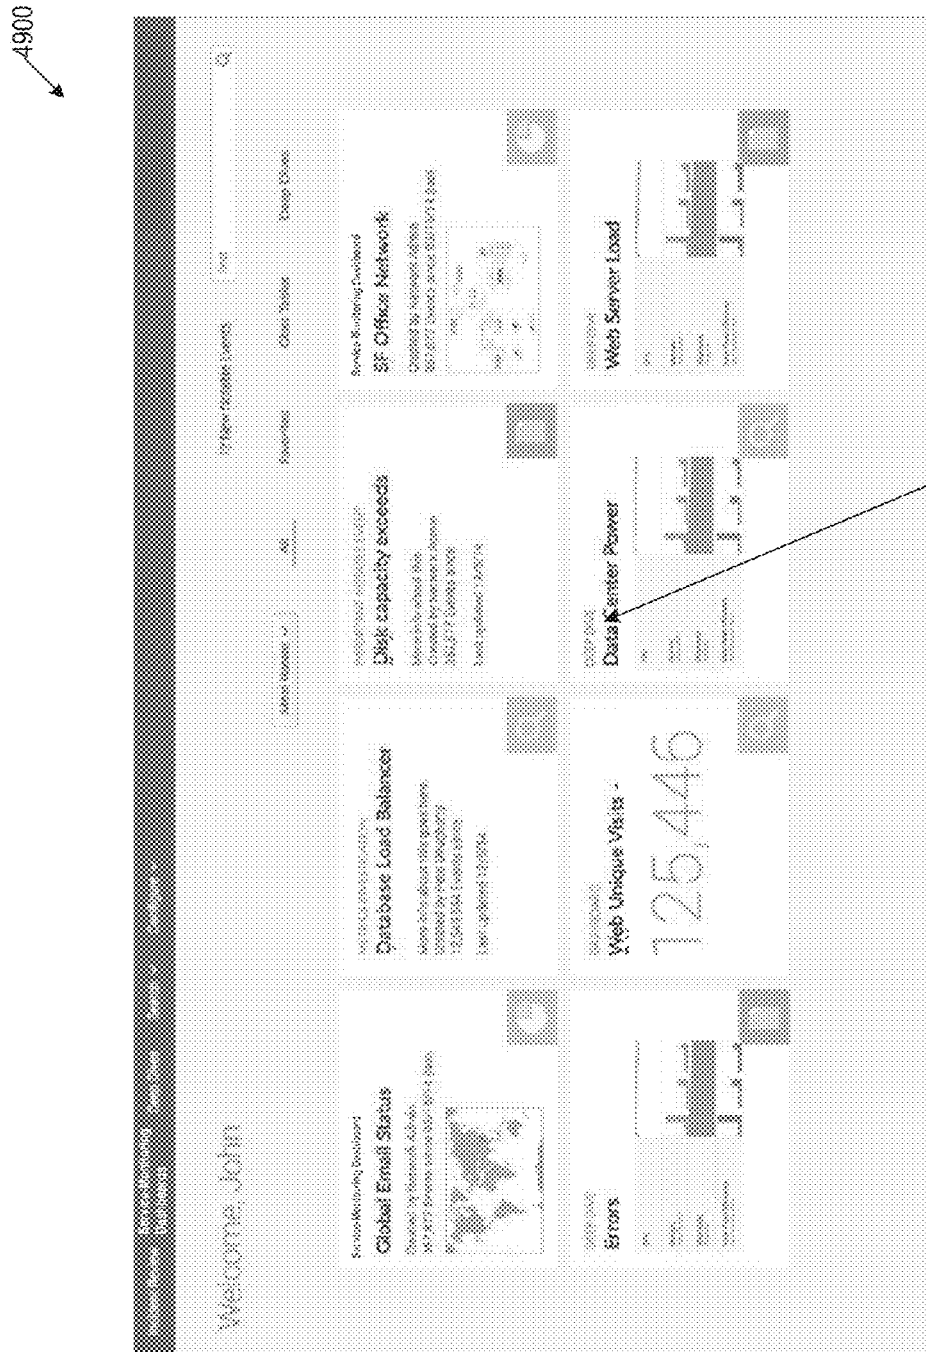


FIG. 49

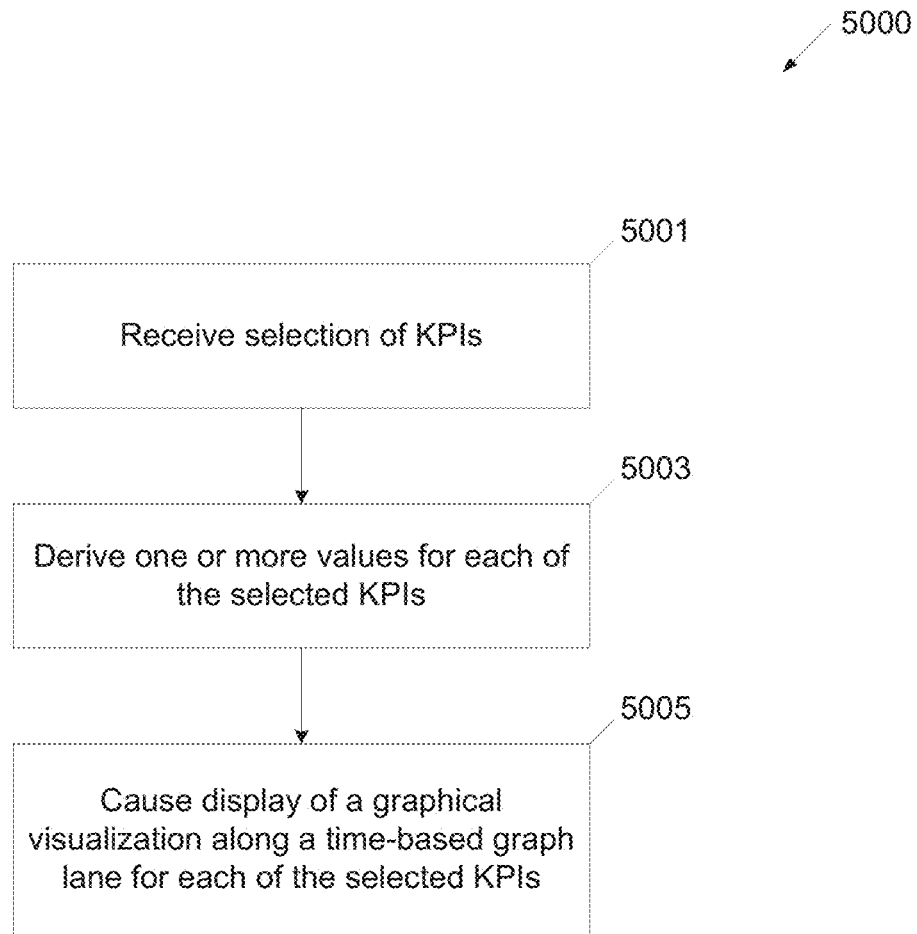
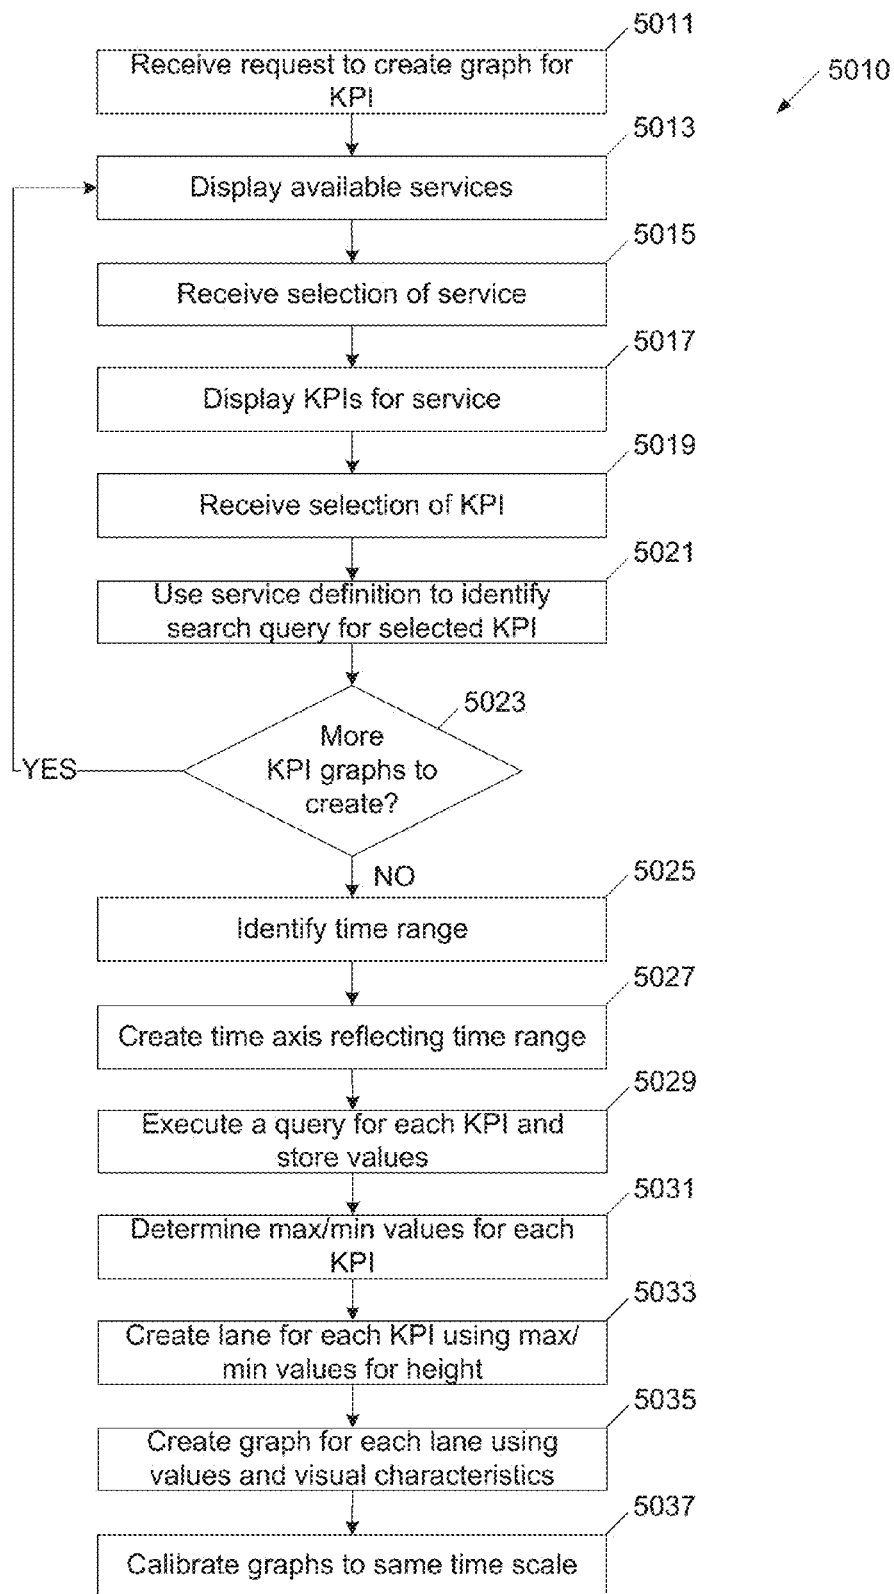


Fig. 50A






Fig. 51

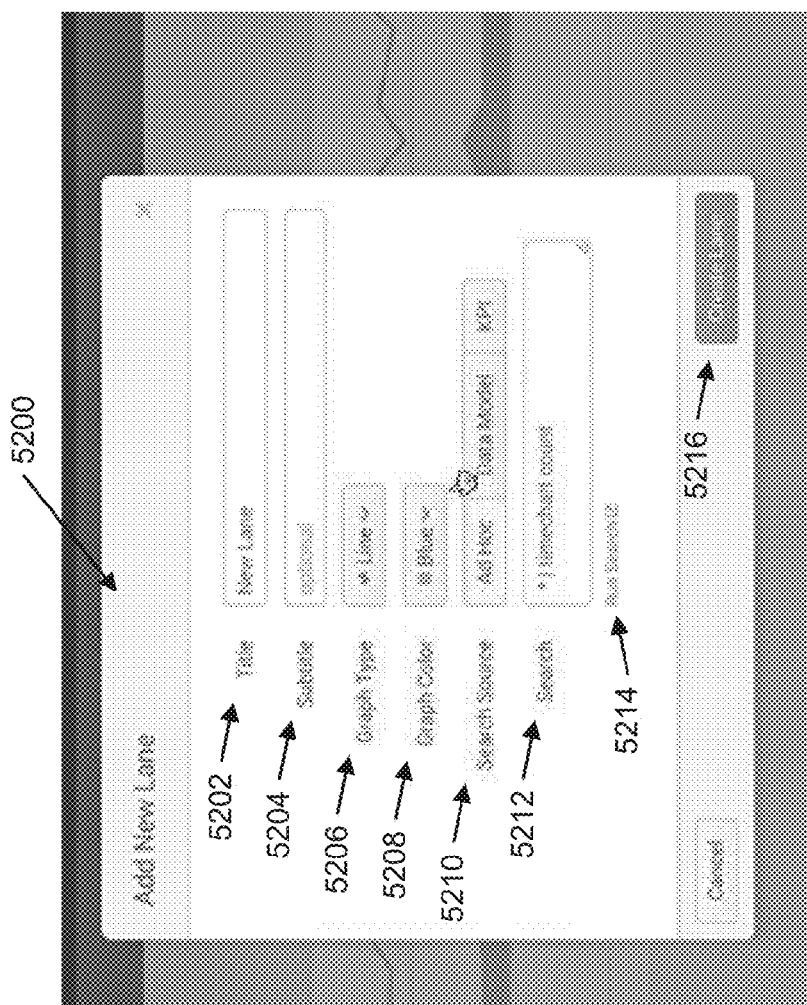


Fig. 52




Fig. 53




Fig. 54






Fig. 55A

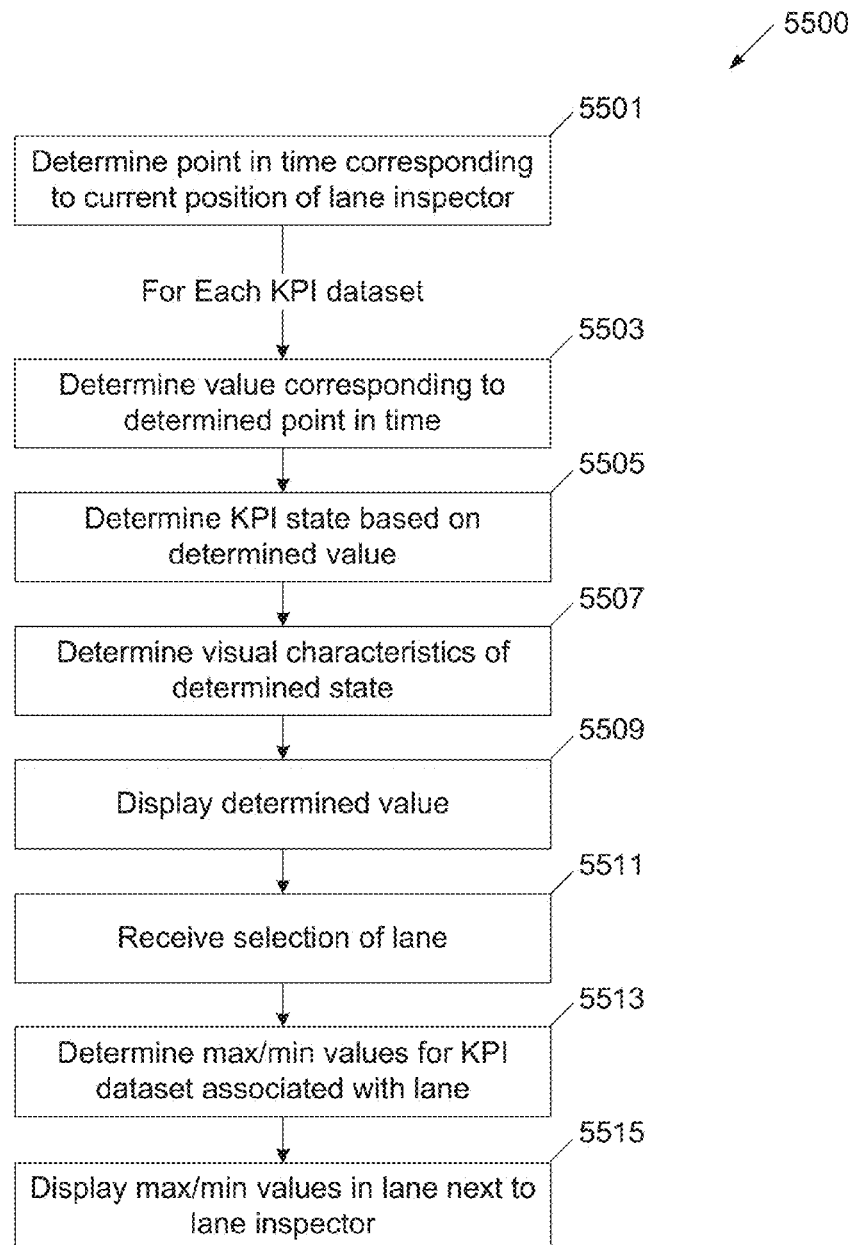


Fig. 55B




Fig. 56

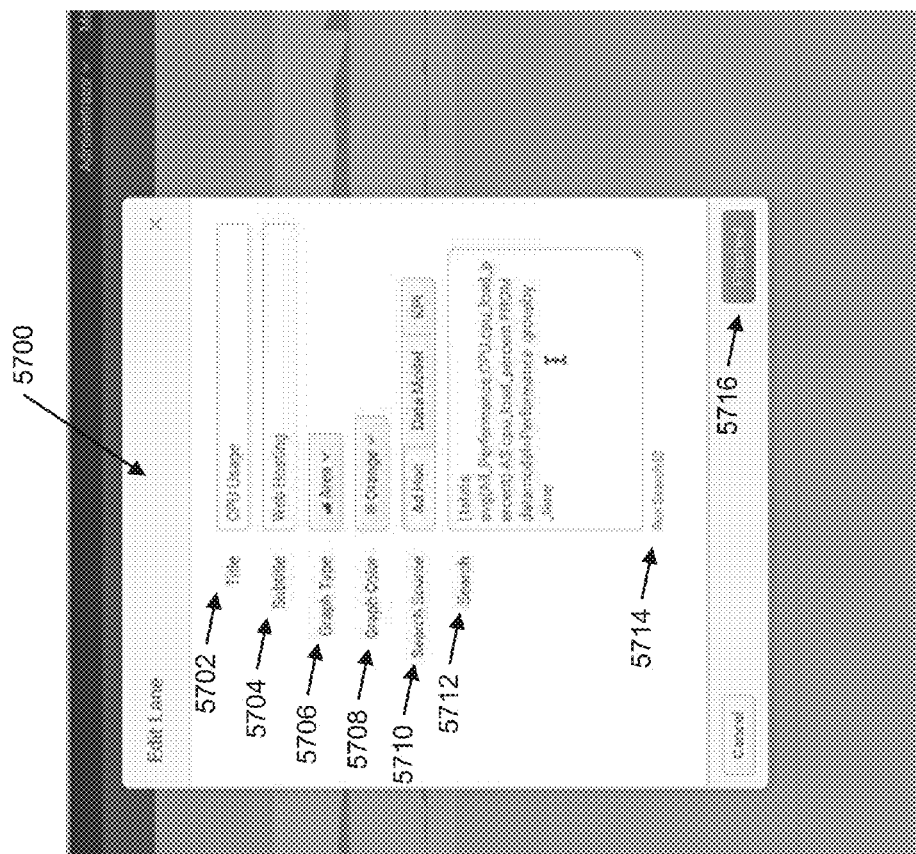


Fig. 57

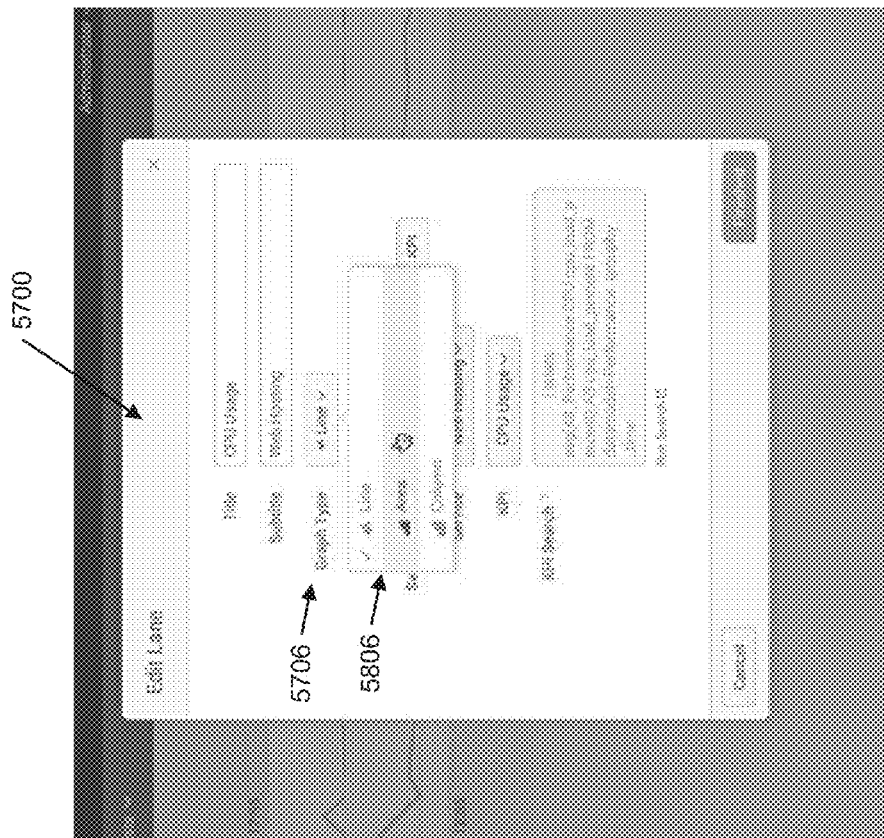


Fig. 58

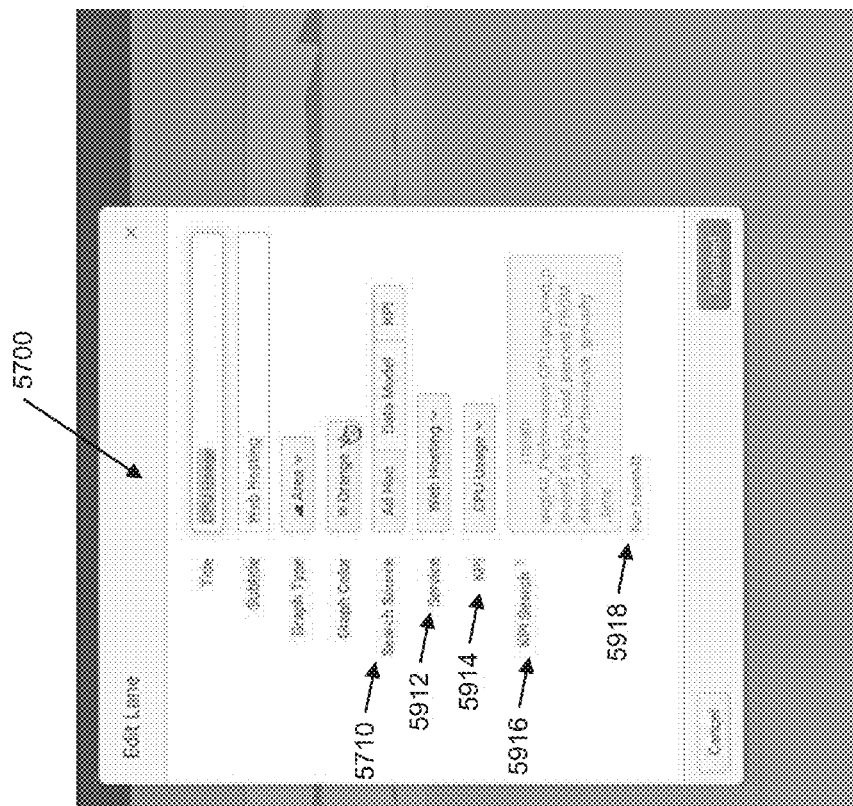


Fig. 59

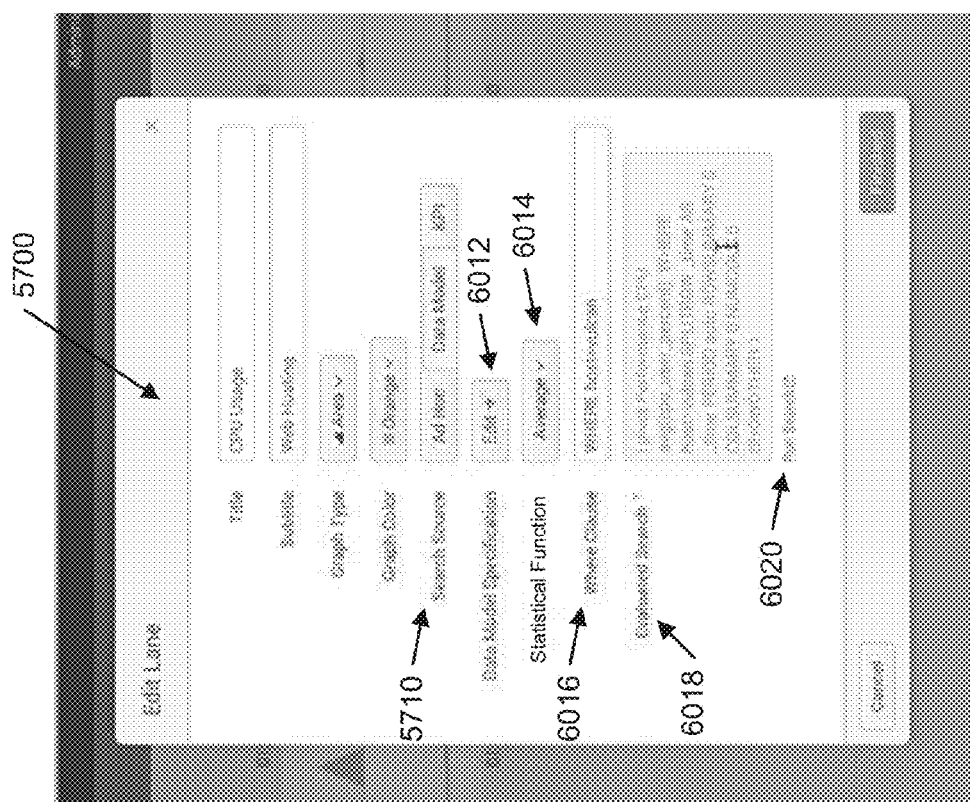


Fig. 60




Fig. 61



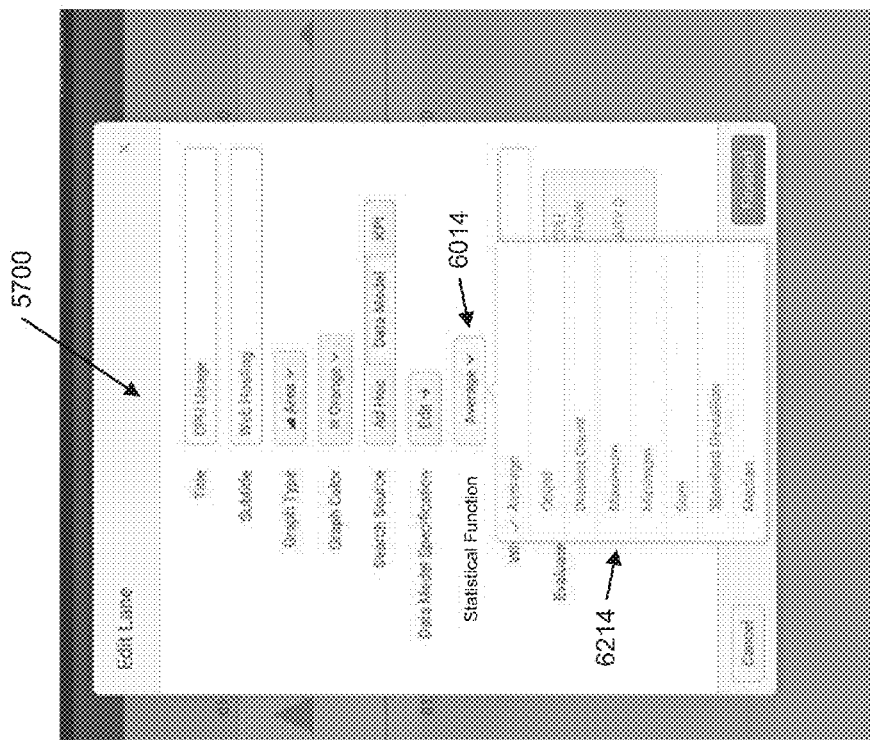


Fig. 62

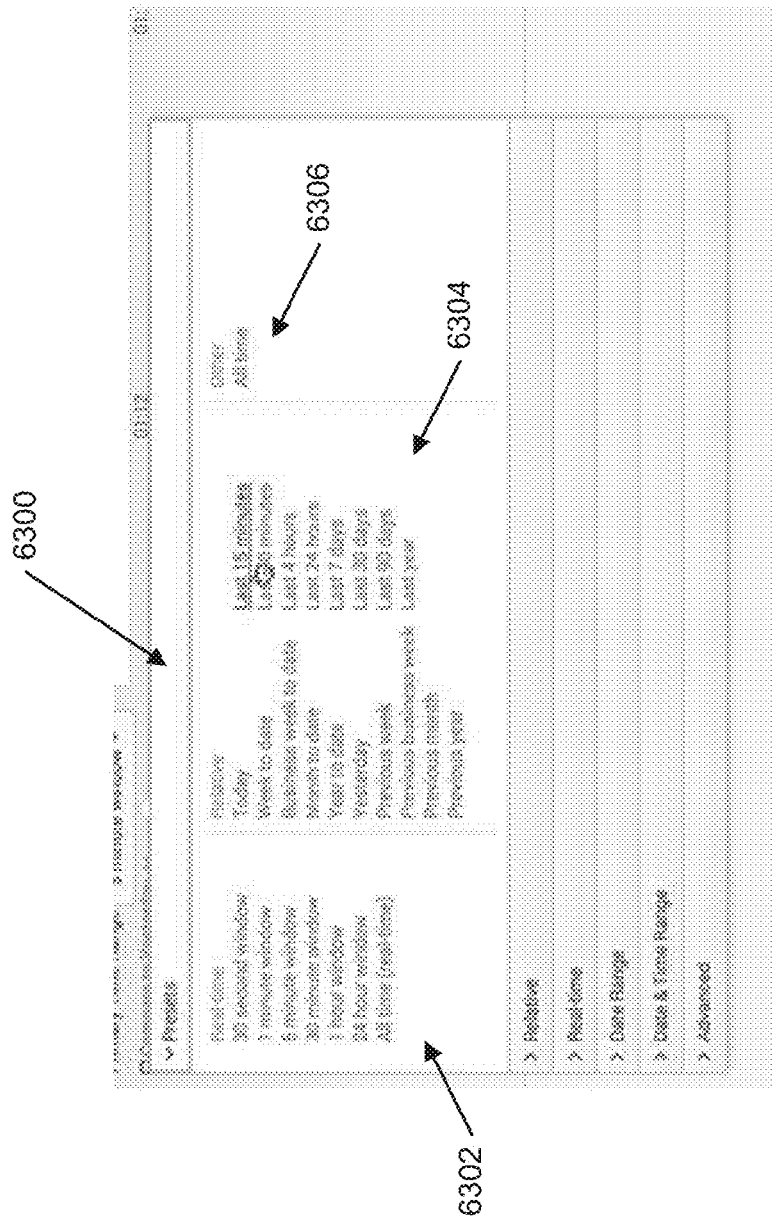


Fig. 63




Fig. 64A

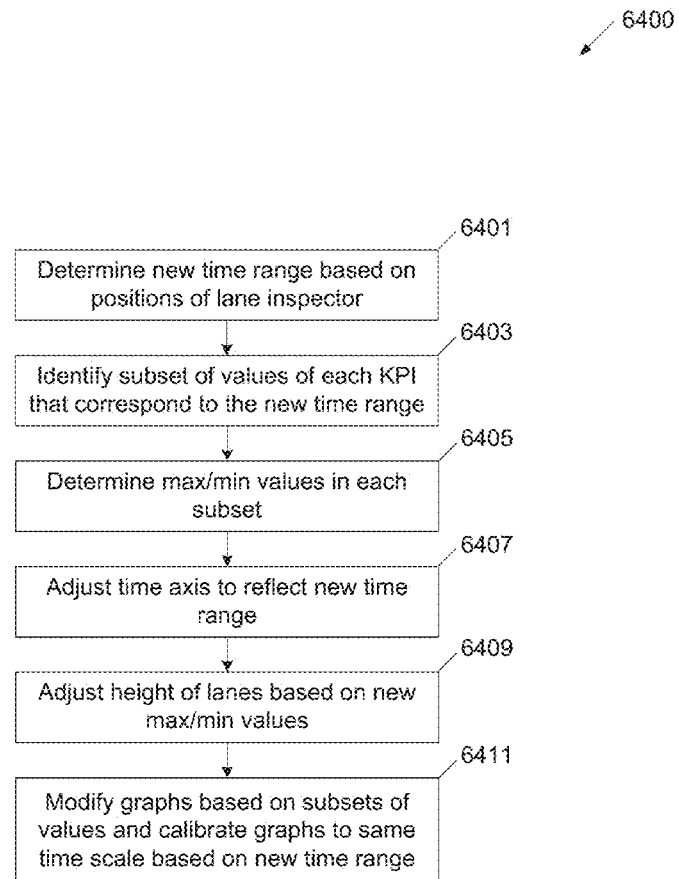


Fig. 64B




Fig. 65




Fig. 66




Fig. 67




Fig. 68






Fig. 69


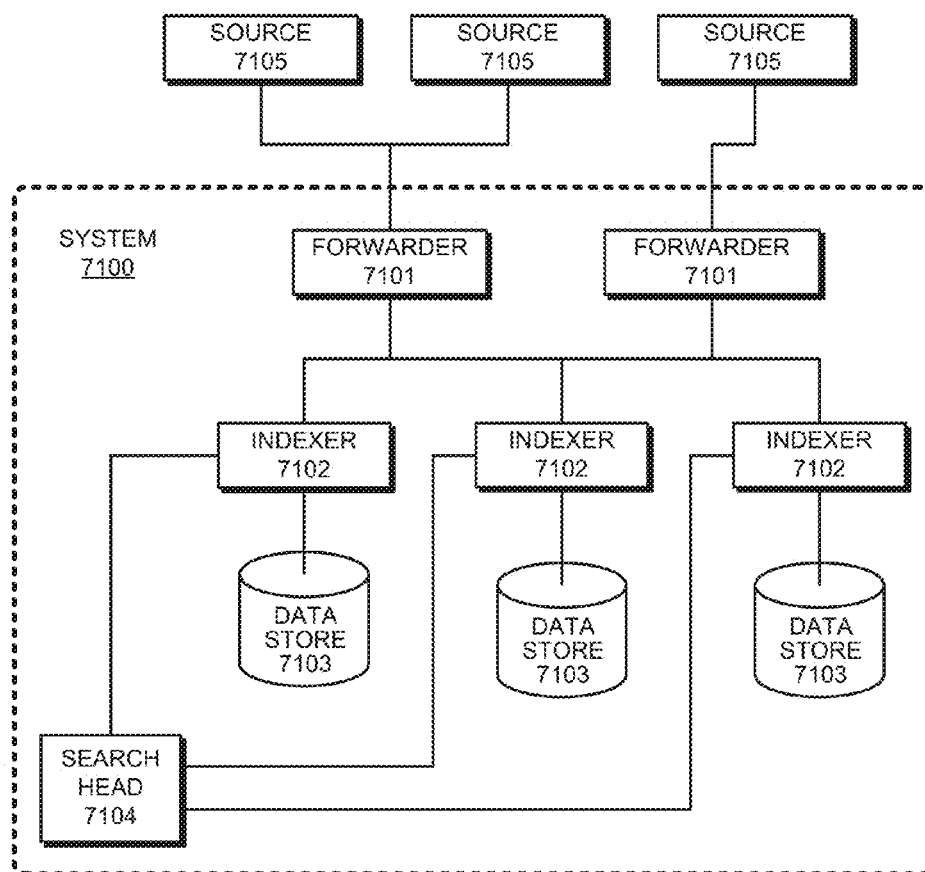
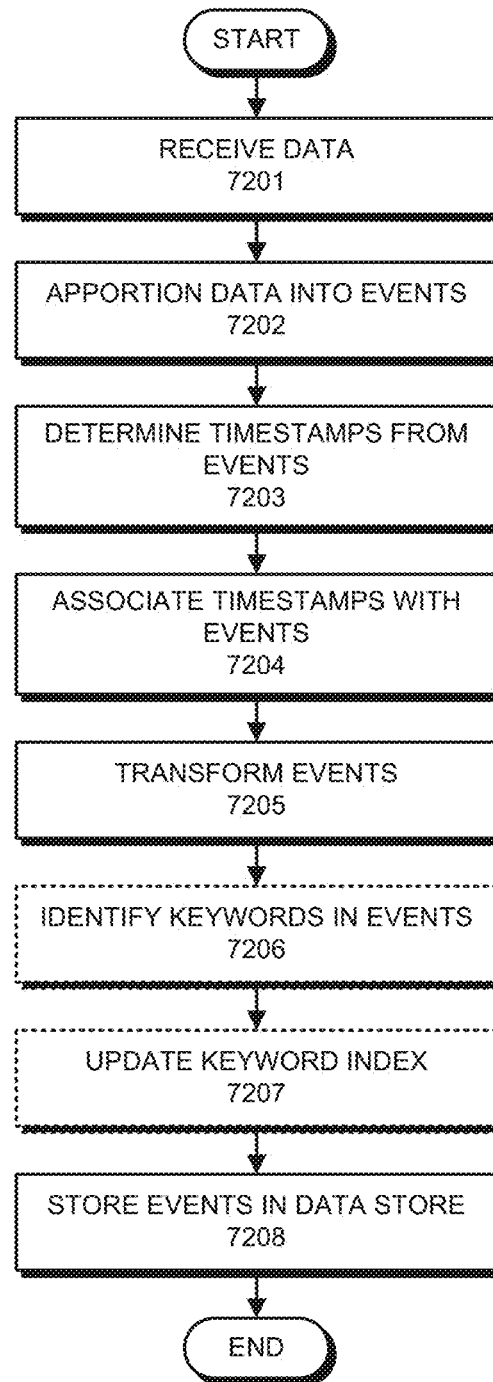

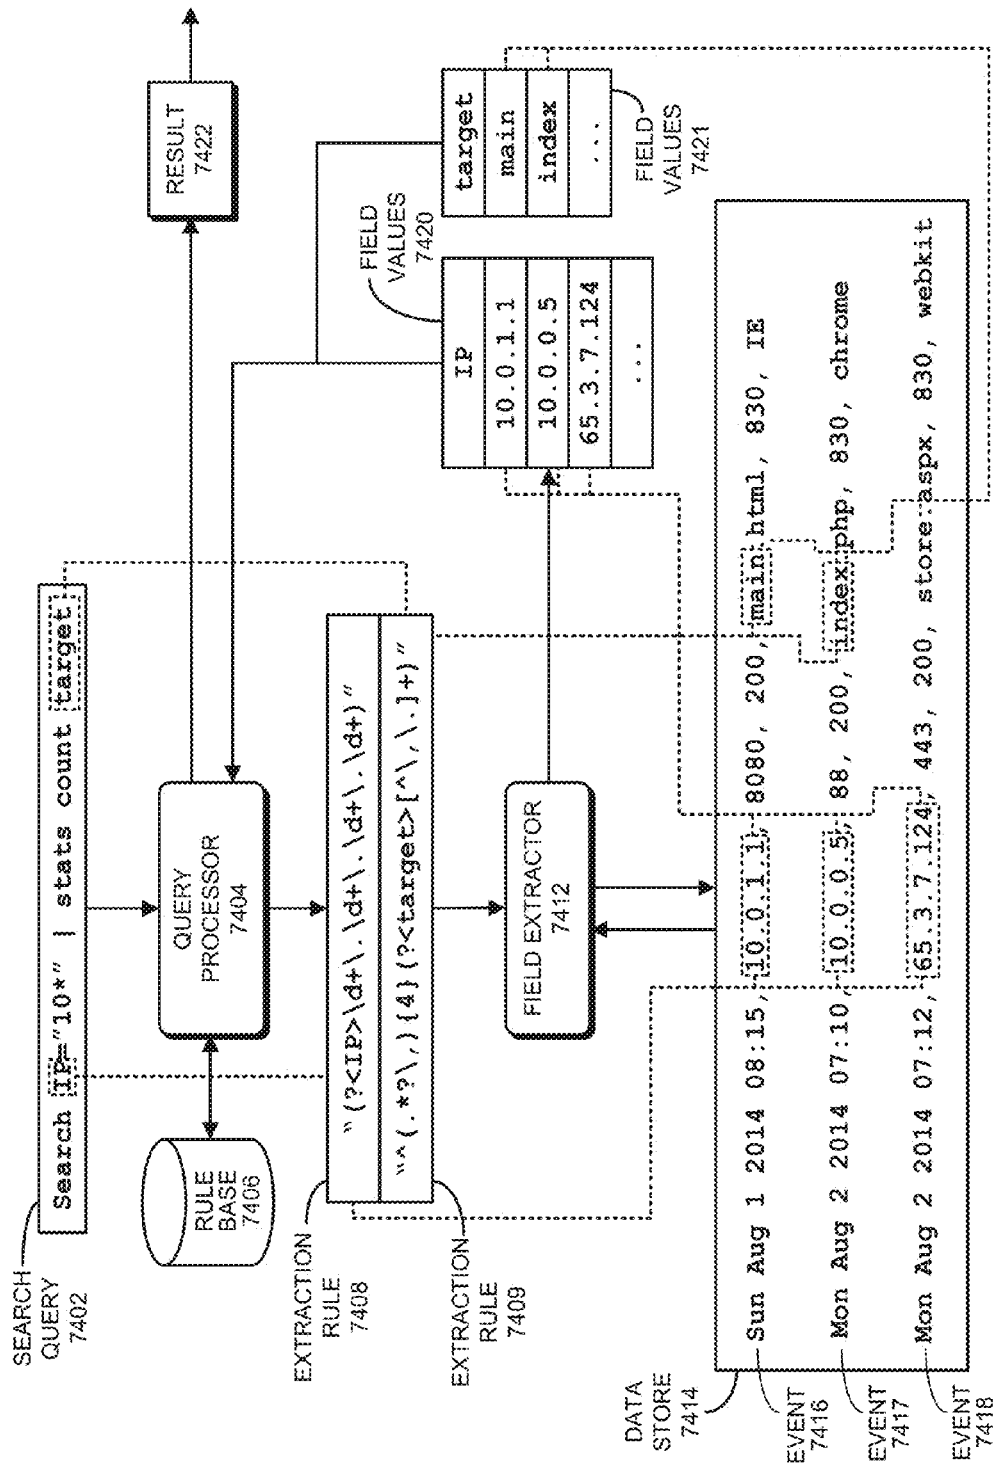


Fig. 70

**FIG. 71**

**FIG. 72**

**FIG. 73**



FG-74A




FIG. 74B

7440

7442

7444

Buttercup Games

Purchases Requests

Purchases Requests

7446

7448

7450

FIG. 74C



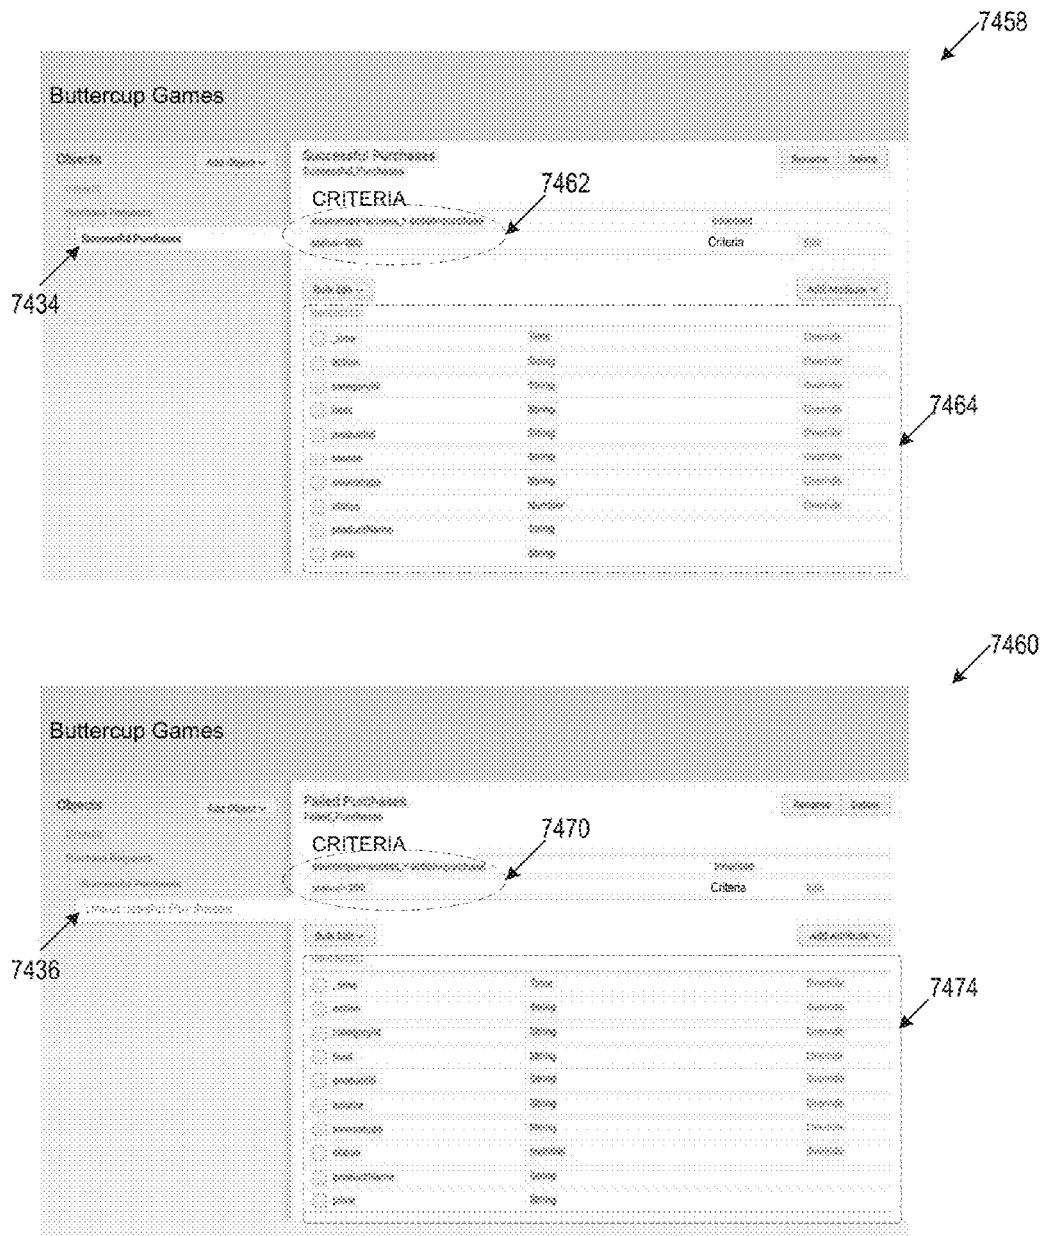


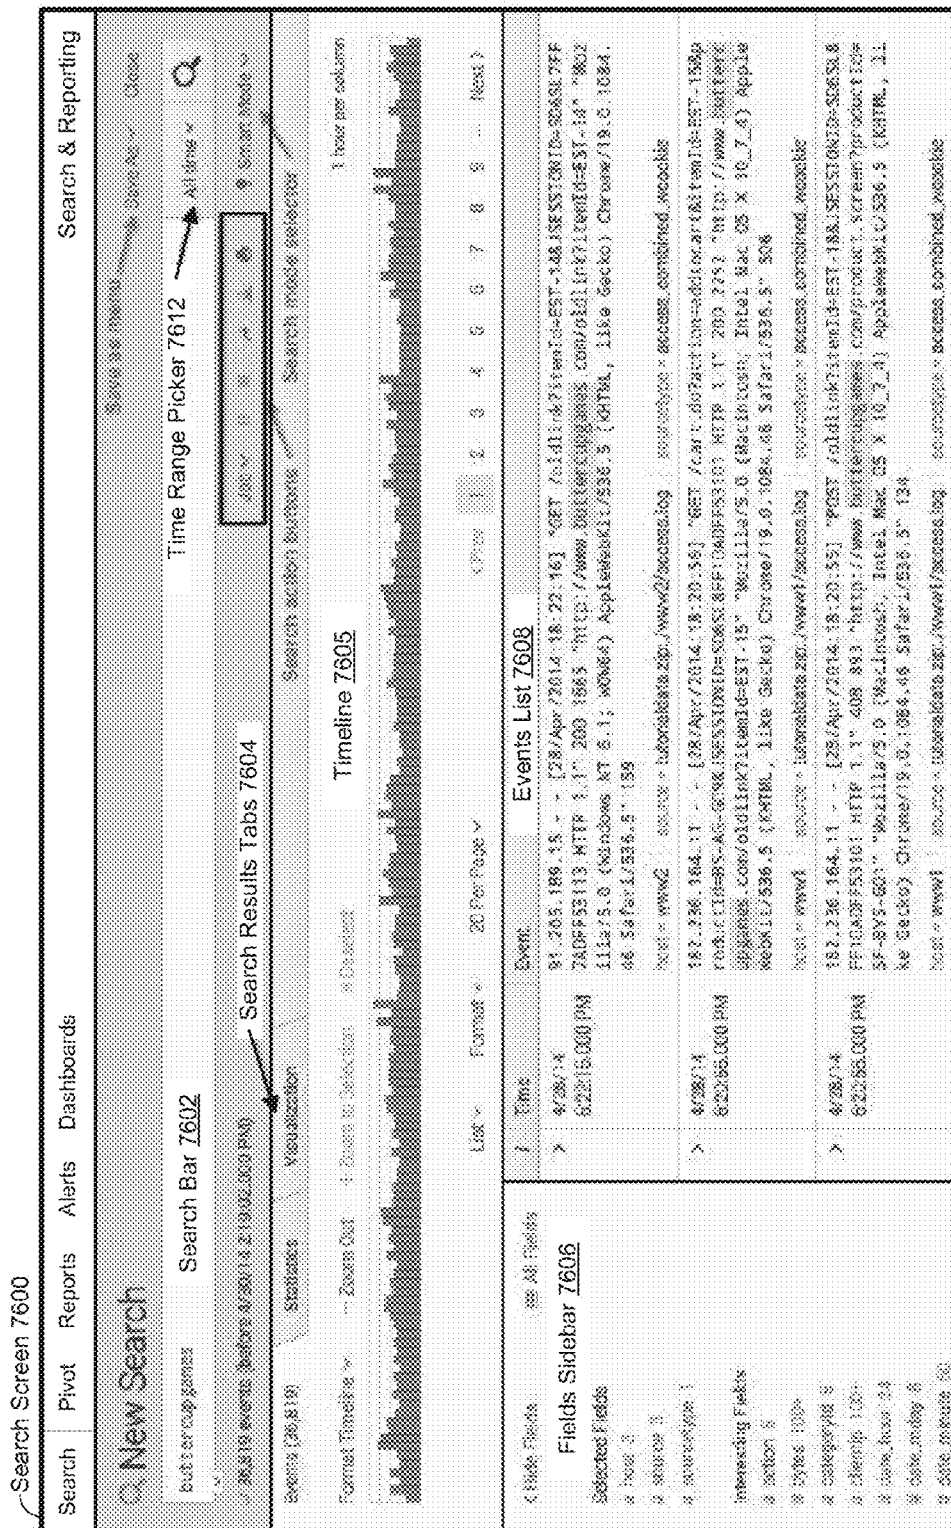
FIG. 74D

Original Search: 7501  
search "error | stats count BY host

Sent to peers: 7502  
search "error | prestats count BY host (map)

Executed by search head: 7503  
Merge prestats results received from peers (reduce)

**FIG. 75**



Data Summary			
Hosts (3)	Sources (3)	SourceTypes (3)	
filter			
Host	id	Count	Last Update
mailov	id	9,829	4/29/14 1:32:47 PM
vendor_data	id	36,244	4/29/14 1:32:46 PM
www1	id	24,221	4/29/14 1:32:44 PM
www2	id	22,568	4/29/14 1:32:47 PM
www3	id	22,375	4/29/14 1:32:46 PM

FIG. 76B

KEY INDICATORS VIEW 7700


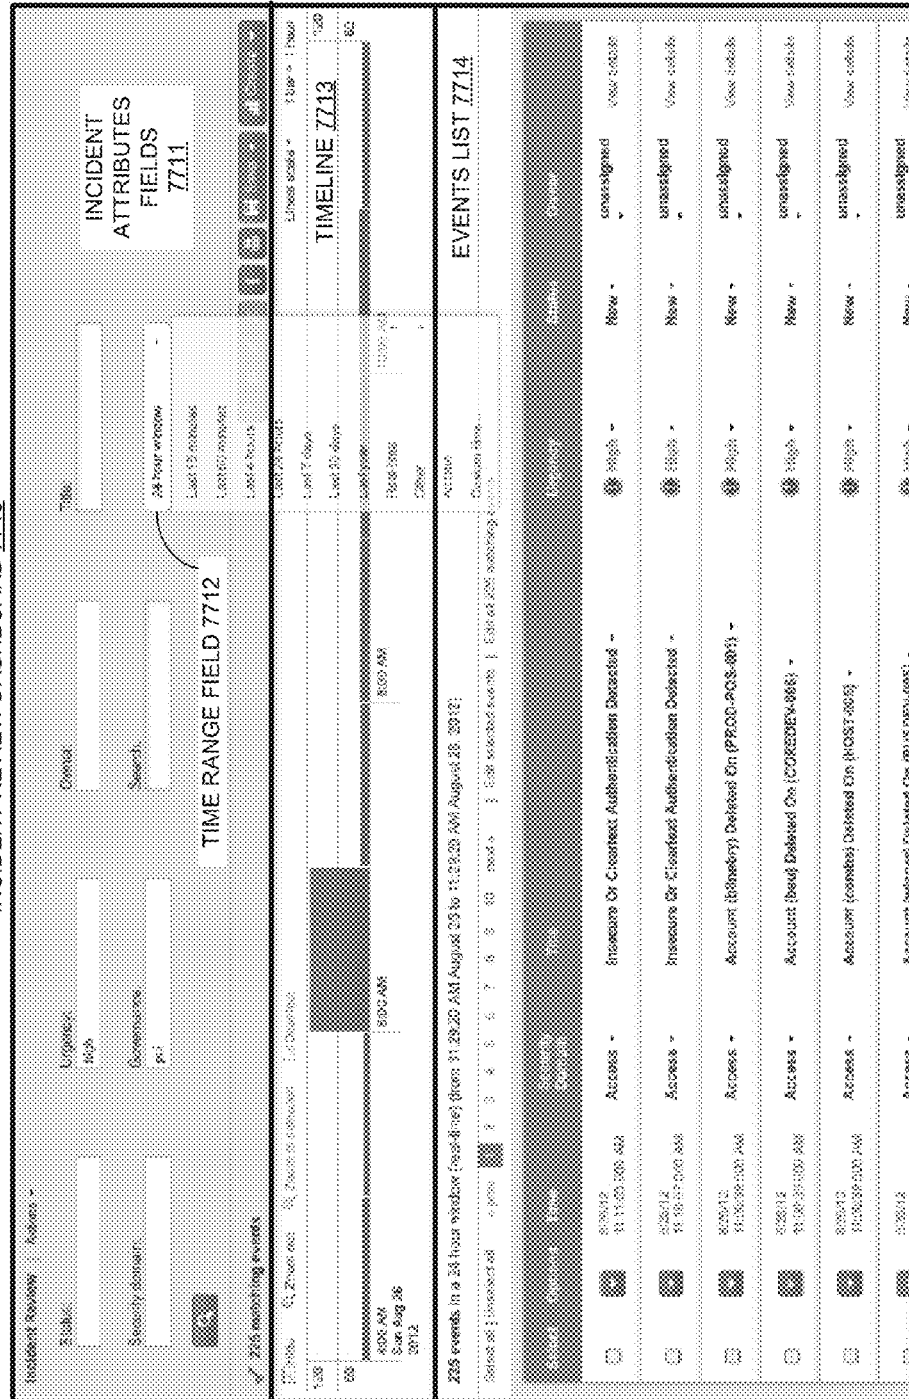


FIG. 77A



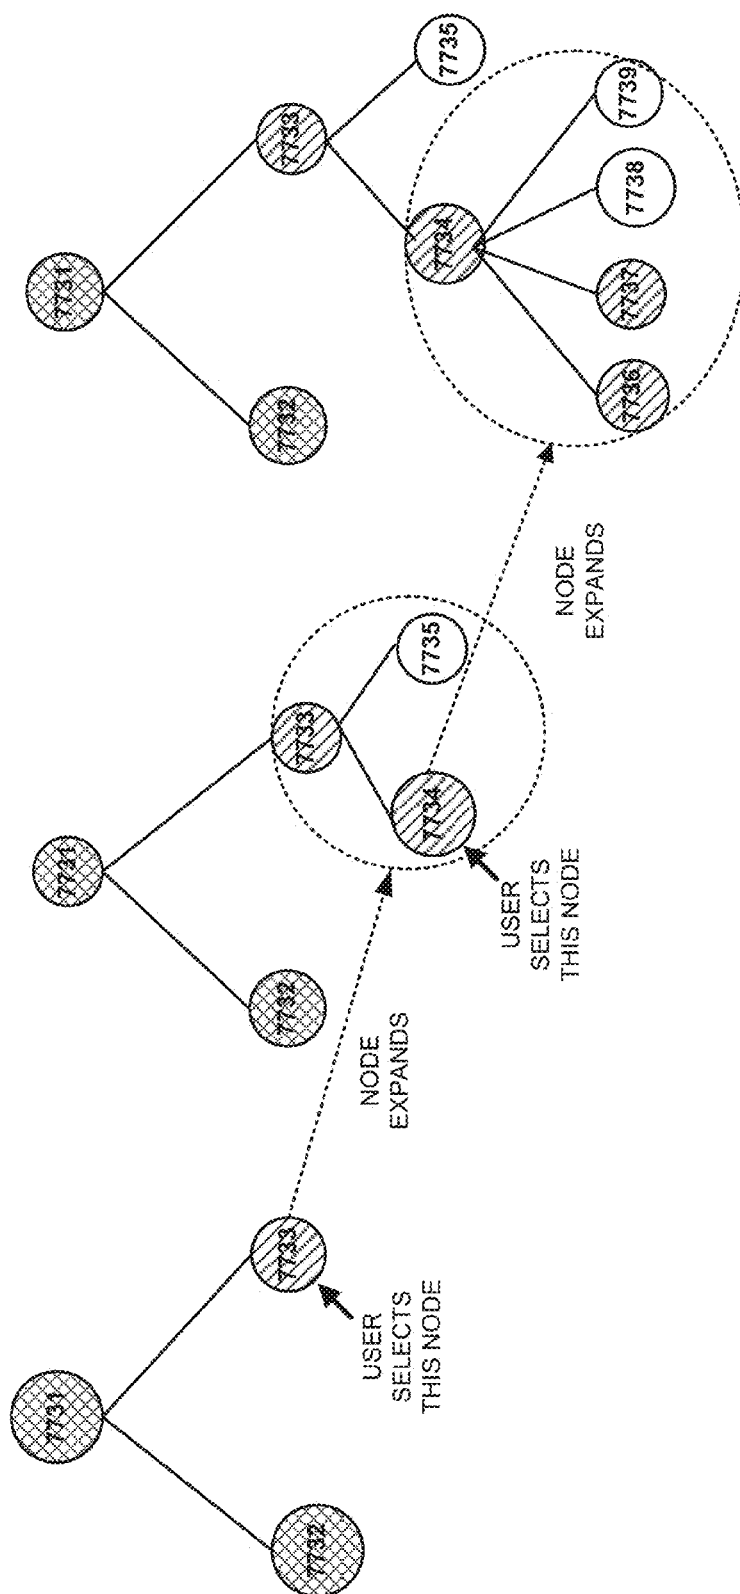



FIG. 77C





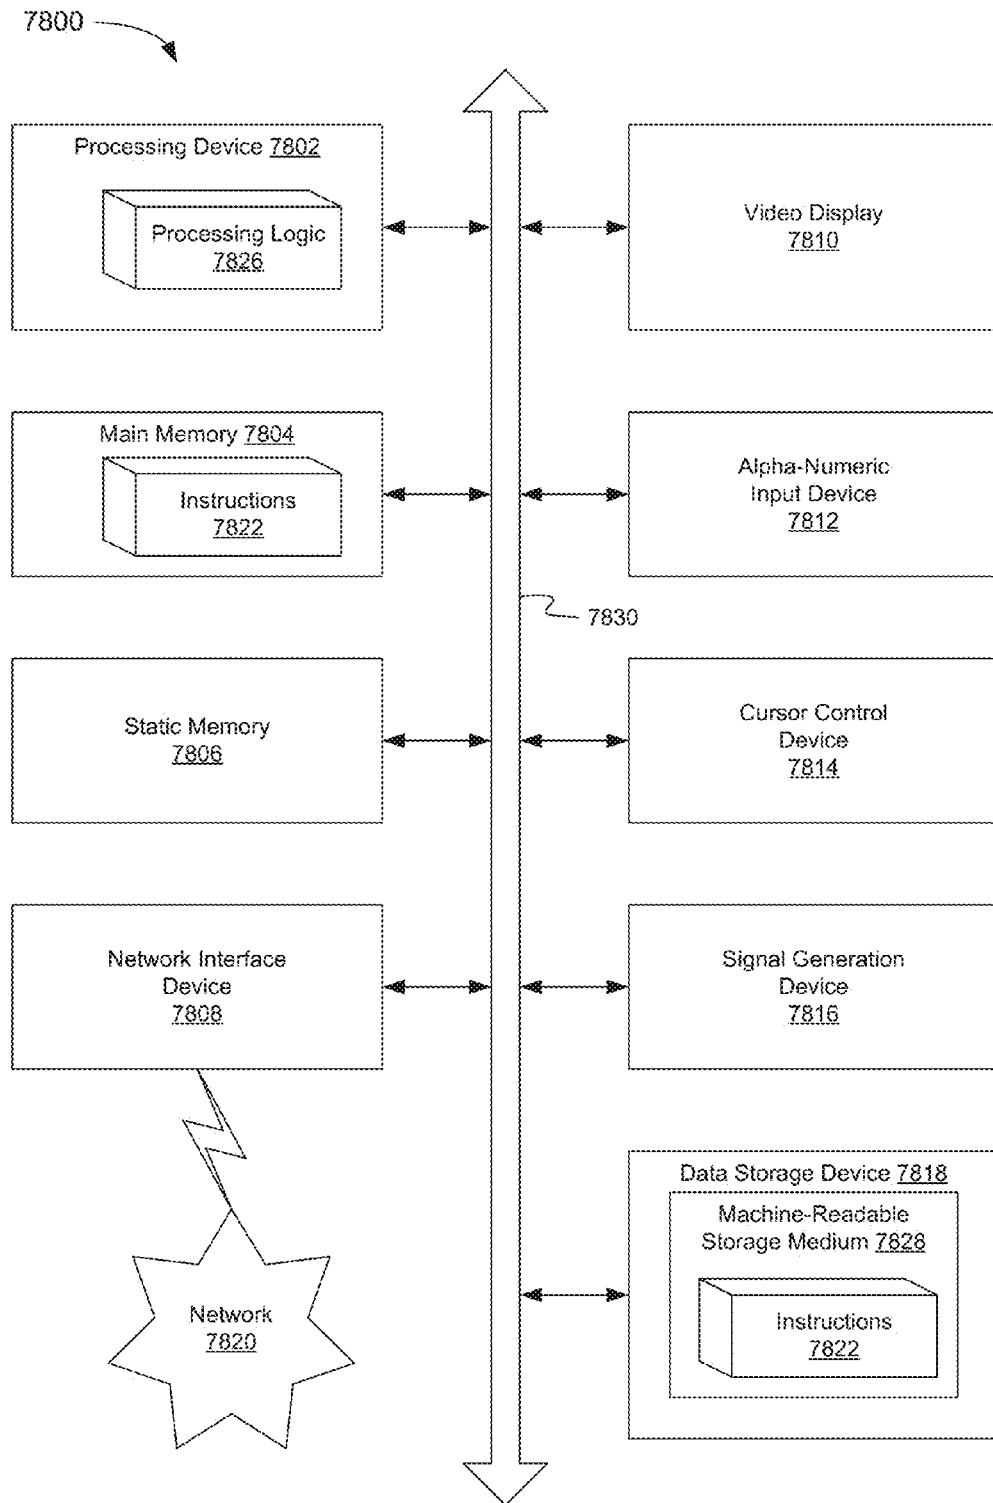


Fig. 78

1

# DEFINING A GRAPHICAL VISUALIZATION ALONG A TIME-BASED GRAPH LANE USING KEY PERFORMANCE INDICATORS DERIVED FROM MACHINE DATA

## RELATED APPLICATION

This application is related to and claims the benefit of U.S. Provisional Patent Application No. 62/062,104 filed Oct. 9, 2014, which is hereby incorporated by reference herein.

## TECHNICAL FIELD

The present disclosure relates to monitoring services and, more particularly, to monitoring service-level performance using key performance indicators derived from machine data.

## BACKGROUND

Modern data centers often comprise thousands of hosts that operate collectively to service requests from even larger numbers of remote clients. During operation, components of these data centers can produce significant volumes of machine-generated data. The unstructured nature of much of this data has made it challenging to perform indexing and searching operations because of the difficulty of applying semantic meaning to unstructured data. As the number of hosts and clients associated with a data center continues to grow, processing large volumes of machine-generated data in an intelligent manner and effectively presenting the results of such processing continues to be a priority.

## BRIEF DESCRIPTION OF THE DRAWINGS

The present disclosure will be understood more fully from the detailed description given below and from the accompanying drawings of various implementations of the disclosure.

FIG. 1 illustrates a block diagram of an example of entities providing a service, in accordance with one or more implementations of the present disclosure.

FIG. 2 is a block diagram of one implementation of a service monitoring system, in accordance with one or more implementations of the present disclosure.

FIG. 3 is a block diagram illustrating an entity definition for an entity, in accordance with one or more implementations of the present disclosure.

FIG. 4 is a block diagram illustrating a service definition that relates one or more entities with a service, in accordance with one or more implementations of the present disclosure.

FIG. 5 is a flow diagram of an implementation of a method for creating one or more key performance indicators for a service, in accordance with one or more implementations of the present disclosure.

FIG. 6 is a flow diagram of an implementation of a method for creating an entity definition for an entity, in accordance with one or more implementations of the present disclosure.

FIG. 7 illustrates an example of a graphical user interface (GUI) for creating and/or editing entity definition(s) and/or service definition(s), in accordance with one or more implementations of the present disclosure.

FIG. 8 illustrates an example of a GUI for creating and/or editing entity definitions, in accordance with one or more implementations of the present disclosure.

FIG. 9A illustrates an example of a GUI for creating an entity definition, in accordance with one or more implementations of the present disclosure.

2

FIG. 9B illustrates an example of input received via GUI for creating an entity definition, in accordance with one or more implementations of the present disclosure.

FIG. 10 illustrates an example of a GUI for creating and/or editing entity definitions, in accordance with one or more implementations of the present disclosure.

FIG. 11 is a flow diagram of an implementation of a method for creating a service definition for a service, in accordance with one or more implementations of the present disclosure.

FIG. 12 illustrates an example of a GUI for creating and/or editing service definitions, in accordance with one or more implementations of the present disclosure.

FIG. 13 illustrates an example of a GUI for identifying a service for a service definition, in accordance with one or more implementations of the present disclosure.

FIG. 14 illustrates an example of a GUI for creating a service definition, in accordance with one or more implementations of the present disclosure.

FIG. 15 illustrates an example of a GUI for associating one or more entities with a service by associating one or more entity definitions with a service definition, in accordance with one or more implementations of the present disclosure.

FIG. 16 illustrates an example of a GUI facilitating user input for creating an entity definition, in accordance with one or more implementations of the present disclosure.

FIG. 17 illustrates an example of a GUI indicating one or more entities associated with a service based on input, in accordance with one or more implementations of the present disclosure.

FIG. 18 illustrates an example of a GUI for specifying dependencies for the service, in accordance with one or more implementations of the present disclosure.

FIG. 19 is a flow diagram of an implementation of a method for creating one or more key performance indicators (KPIs) for a service, in accordance with one or more implementations of the present disclosure.

FIG. 20 is a flow diagram of an implementation of a method for creating a search query, in accordance with one or more implementations of the present disclosure.

FIG. 21 illustrates an example of a GUI for creating a KPI for a service, in accordance with one or more implementations of the present disclosure.

FIG. 22 illustrates an example of a GUI for creating a KPI for a service, in accordance with one or more implementations of the present disclosure.

FIG. 23 illustrates an example of a GUI for receiving input of search processing language for defining a search query for a KPI for a service, in accordance with one or more implementations of the present disclosure.

FIG. 24 illustrates an example of a GUI for defining a search query for a KPI using a data model, in accordance with one or more implementations of the present disclosure.

FIG. 25 illustrates an example of a GUI for facilitating user input for selecting a data model and an object of the data model to use for the search query, in accordance with one or more implementations of the present disclosure.

FIG. 26 illustrates an example of a GUI for displaying a selected statistic, in accordance with one or more implementations of the present disclosure.

FIG. 27 illustrates an example of a GUI for editing which entity definitions to use for the KPI, in accordance with one or more implementations of the present disclosure.

FIG. 28 is a flow diagram of an implementation of a method for defining one or more thresholds for a KPI, in accordance with one or more implementations of the present disclosure.

FIGS. 29A-B, illustrate examples of a graphical interface enabling a user to set a threshold for the KPI, in accordance with one or more implementations of the present disclosure.

FIG. 30 illustrates an example GUI for enabling a user to set one or more thresholds for the KPI, in accordance with one or more implementations of the present disclosure.

FIG. 31A-C illustrate example GUIs for defining thresholds for a KPI, in accordance with one or more implementations of the present disclosure.

FIG. 32 is a flow diagram of an implementation of a method for calculating an aggregate KPI score for a service based on the KPIs for the service, in accordance with one or more implementations of the present disclosure.

FIG. 33A illustrates an example GUI 3300 for assigning a frequency of monitoring to a KPI based on user input, in accordance with one or more implementations of the present disclosure.

FIG. 33B illustrates an example GUI for defining threshold settings, including state ratings, for a KPI, in accordance with one or more implementations of the present disclosure.

FIG. 34 is a flow diagram of an implementation of a method for calculating a value for an aggregate KPI for the service, in accordance with one or more implementations of the present disclosure.

FIG. 35 is a flow diagram of an implementation of a method for creating a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure.

FIG. 36A illustrates an example GUI for creating and/or editing a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure.

FIG. 36B illustrates an example GUI for a dashboard-creation graphical interface for creating a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure.

FIG. 37 illustrates an example GUI for a dashboard-creation graphical interface including a user selected background image, in accordance with one or more implementations of the present disclosure.

FIG. 38 illustrates an example GUI for displaying of a set of KPIs associated with a selected service, in accordance with one or more implementations of the present disclosure.

FIG. 39 illustrates an example GUI facilitating user input for selecting a location in the dashboard template and style settings for a KPI widget, and displaying the KPI widget in the dashboard template, in accordance with one or more implementations of the present disclosure.

FIG. 40 illustrates an example Noel gauge widget, in accordance with one or more implementations of the present disclosure.

FIG. 41 illustrates an example single value widget, in accordance with one or more implementations of the present disclosure.

FIG. 42 illustrates an example GUI illustrating a search query and a search result for a Noel gauge widget, a single value widget, and a trend indicator widget, in accordance with one or more implementations of the present disclosure.

FIG. 43 illustrates an example GUI portion of a service-monitoring dashboard for facilitating user input specifying a time range to use when executing a search query defining a KPI, in accordance with one or more implementations of the present disclosure.

FIG. 44 illustrates spark line widget, in accordance with one or more implementations of the present disclosure.

FIG. 45 illustrates an example GUI illustrating a search query and search results for a spark line widget, in accordance with one or more implementations of the present disclosure.

FIG. 46 illustrates a trend indicator widget, in accordance with one or more implementations of the present disclosure.

FIG. 47A is a flow diagram of an implementation of a method for creating and causing for display a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure.

FIG. 47B describes an example service-monitoring dashboard GUI, in accordance with one or more implementations of the present disclosure.

FIG. 48 describes an example home page GUI for service-level monitoring, in accordance with one or more implementations of the present disclosure.

FIG. 49 describes an example home page GUI for service-level monitoring, in accordance with one or more implementations of the present disclosure.

FIG. 50A is a flow diagram of an implementation of a method for creating a visual interface displaying graphical visualizations of KPI values along time-based graph lanes, in accordance with one or more implementations of the present disclosure.

FIG. 50B is a flow diagram of an implementation of a method for generating a graphical visualization of KPI values along a time-based graph lane, in accordance with one or more implementations of the present disclosure.

FIG. 51 illustrates an example of a graphical user interface (GUI) for creating a visual interface displaying graphical visualizations of KPI values along time-based graph lanes, in accordance with one or more implementations of the present disclosure.

FIG. 52 illustrates an example of a GUI for adding a graphical visualization of KPI values along a time-based graph lane to a visual interface, in accordance with one or more implementations of the present disclosure.

FIG. 53 illustrates an example of a visual interface with time-based graph lanes for displaying graphical visualizations, in accordance with one or more implementations of the present disclosure.

FIG. 54 illustrates an example of a visual interface displaying graphical visualizations of KPI values along time-based graph lanes, in accordance with one or more implementations of the present disclosure.

FIG. 55A illustrates an example of a visual interface with a user manipulable visual indicator spanning across the time-based graph lanes, in accordance with one or more implementations of the present disclosure.

FIG. 55B is a flow diagram of an implementation of a method for inspecting graphical visualizations of KPI values along a time-based graph lane, in accordance with one or more implementations of the present disclosure.

FIG. 56 illustrates an example of a visual interface displaying graphical visualizations of KPI values along time-based graph lanes with options for editing the graphical visualizations, in accordance with one or more implementations of the present disclosure.

FIG. 57 illustrates an example of a GUI for editing a graphical visualization of KPI values along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure.

FIG. 58 illustrates an example of a GUI for editing a graph style of a graphical visualization of KPI values along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure.

FIG. 59 illustrates an example of a GUI for selecting the KPI corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure.

5

FIG. 60 illustrates an example of a GUI for selecting a data model corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure.

FIG. 61 illustrates an example of a GUI for selecting a data model corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure.

FIG. 62 illustrates an example of a GUI for editing an aggregation operation for a data model corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure.

FIG. 63 illustrates an example of a GUI for selecting a time range that graphical visualizations along a time-based graph lane in a visual interface should cover, in accordance with one or more implementations of the present disclosure.

FIG. 64A illustrates an example of a visual interface for selecting a subset of a time range that graphical visualizations along a time-based graph lane in a visual interface cover, in accordance with one or more implementations of the present disclosure.

FIG. 64B is a flow diagram of an implementation of a method for enhancing a view of a subset a subset of a time range for a time-based graph lane, in accordance with one or more implementations of the present disclosure.

FIG. 65 illustrates an example of a visual interface displaying graphical visualizations of KPI values along time-based graph lanes for a selected subset of a time range, in accordance with one or more implementations of the present disclosure.

FIG. 66 illustrates an example of a visual interface displaying twin graphical visualizations of KPI values along time-based graph lanes for different periods of time, in accordance with one or more implementations of the present disclosure.

FIG. 67 illustrates an example of a visual interface with a user manipulable visual indicator spanning across twin graphical visualizations of KPI values along time-based graph lanes for different periods of time, in accordance with one or more implementations of the present disclosure.

FIG. 68 illustrates an example of a visual interface displaying a graph lane with inventory information for a service or entities reflected by KPI values, in accordance with one or more implementations of the present disclosure.

FIG. 69 illustrates an example of a visual interface displaying a graph lane with notable events occurring during a timer period covered by graphical visualization of KPI values, in accordance with one or more implementations of the present disclosure.

FIG. 70 illustrates an example of a visual interface displaying a graph lane with notable events occurring during a timer period covered by graphical visualization of KPI values, in accordance with one or more implementations of the present disclosure.

FIG. 71 presents a block diagram of an event-processing system in accordance with one or more implementations of the present disclosure.

FIG. 72 presents a flowchart illustrating how indexers process, index, and store data received from forwarders in accordance with one or more implementations of the present disclosure.

FIG. 73 presents a flowchart illustrating how a search head and indexers perform a search query in accordance with one or more implementations of the present disclosure.

6

FIG. 74A presents a block diagram of a system for processing search requests that uses extraction rules for field values in accordance with one or more implementations of the present disclosure.

FIG. 74B illustrates an example data model structure, in accordance with some implementations of the present disclosure.

FIG. 74C illustrates an example definition of a root object of a data model, in accordance with some implementations.

FIG. 74D illustrates example definitions and of child objects, in accordance with some implementations.

FIG. 75 illustrates an exemplary search query received from a client and executed by search peers in accordance with one or more implementations of the present disclosure.

FIG. 76A illustrates a search screen in accordance with one or more implementations of the present disclosure.

FIG. 76B illustrates a data summary dialog that enables a user to select various data sources in accordance with one or more implementations of the present disclosure.

FIG. 77A illustrates a key indicators view in accordance with one or more implementations of the present disclosure.

FIG. 77B illustrates an incident review dashboard in accordance with one or more implementations of the present disclosure.

FIG. 77C illustrates a proactive monitoring tree in accordance with one or more implementations of the present disclosure.

FIG. 77D illustrates a screen displaying both log data and performance data in accordance with one or more implementations of the present disclosure.

FIG. 78 depicts a block diagram of an example computing device operating in accordance with one or more implementations of the present disclosure.

## DETAILED DESCRIPTION

### Overview

The present disclosure is directed to monitoring performance of a system at a service level using key performance indicators derived from machine data. Implementations of the present disclosure provide users with insight to the performance of monitored services, such as, services pertaining to an information technology (IT) environment. For example, one or more users may wish to monitor the performance of a web hosting service, which provides hosted web content to end users via network.

A service can be provided by one or more entities. An entity that provides a service can be associated with machine data. As described in greater detail below, the machine data pertaining to a particular entity may use different formats and/or different aliases for the entity.

Implementations of the present disclosure are described for normalizing the different aliases and/or formats of machine data pertaining to the same entity. In particular, an entity definition can be created for a respective entity. The entity definition can normalize various machine data pertaining to a particular entity, thus simplifying the use of heterogeneous machine data for monitoring a service.

Implementations of the present disclosure are described for specifying which entities, and thus, which heterogeneous machine data, to use for monitoring a service. In one implementation, a service definition is created for a service that is to be monitored. The service definition specifies one or more entity definitions, where each entity definition corresponds to a respective entity providing the service. The service definition provides users with flexibility in associating entities with services. The service definition further provides users with

the ability to define relationships between entities and services at the machine data level. Implementations of the present disclosure enable end-users to monitor services from a top-down perspective and can provide rich visualization to troubleshoot any service-related issues. Implementations of the present disclosure enable end-users to understand an environment (e.g., IT environment) and the services in the environment. For example, end-users can understand and monitor services at a business service level, application tier level, etc.

Implementations of the present disclosure are described for monitoring a service at a granular level. For example, one or more aspects of a service can be monitored using one or more key performance indicators for the service. A performance indicator or key performance indicator (KPI) is a type of performance measurement. For example, users may wish to monitor the CPU (central processing unit) usage of a web hosting service, the memory usage of the web hosting service, and the request response time for the web hosting service. In one implementation, a separate KPI can be created for each of these aspects of the service that indicates how the corresponding aspect is performing.

Implementations of the present disclosure give users freedom to decide which aspects to monitor for a service and which heterogeneous machine data to use for a particular KPI. In particular, one or more KPIs can be created for a service. Each KPI can be defined by a search query that produces a value derived from the machine data identified in the entity definitions specified in the service definition. Each value can be indicative of how a particular aspect of the service is performing at a point in time or during a period of time. Implementations of the present disclosure enable users to decide what value should be produced by the search query defining the KPI. For example, a user may wish that the request response time be monitored as the average response time over a period of time.

Implementations of the present disclosure are described for customizing various states that a KPI can be in. For example, a user may define a Normal state, a Warning state, and a Critical state for a KPI, and the value produced by the search query of the KPI can indicate the current state of the KPI. In one implementation, one or more thresholds are created for each KPI. Each threshold defines an end of a range of values that represent a particular state of the KPI. A graphical interface can be provided to facilitate user input for creating one or more thresholds for each KPI, naming the states for the KPI, and associating a visual indicator (e.g., color, pattern) to represent a respective state.

Implementations of the present disclosure are described for monitoring a service at a more abstract level, as well. In particular, an aggregate KPI can be configured and calculated for a service to represent the overall health of a service. For example, a service may have 10 KPIs, each monitoring a various aspect of the service. The service may have 7 KPIs in a Normal state, 2 KPIs in a Warning state, and 1 KPI in a Critical state. The aggregate KPI can be a value representative of the overall performance of the service based on the values for the individual KPIs. Implementations of the present disclosure allow individual KPIs of a service to be weighted in terms of how important a particular KPI is to the service relative to the other KPIs in the service, thus giving users control of how to represent the overall performance of a service and control in providing a more accurate representation of the performance of the service. In addition, specific actions can be defined that are to be taken when the aggregate KPI indicating the overall health of a service, for example, exceeds a particular threshold.

Implementations of the present disclosure are described for creating notable events and/or alarms via distribution thresholding. In one implementation, a correlation search is created and used to generate notable event(s) and/or alarm(s). A correlation search can be created to determine the status of a set of KPIs for a service over a defined window of time. A correlation search represents a search query that has a triggering condition and one or more actions that correspond to the trigger condition. Thresholds can be set on the distribution of the state of each individual KPI and if the distribution thresholds are exceeded then an alert/alarm can be generated.

Implementations of the present disclosure are described for providing a service-monitoring dashboard that displays one or more KPI widgets. Each KPI widget can provide a numerical or graphical representation of one or more values for a corresponding KPI or service health score (aggregate KPI for a service) indicating how a service or an aspect of a service is performing at one or more points in time. Users can be provided with the ability to design and draw the service-monitoring dashboard and to customize each of the KPI widgets. A dashboard-creation graphical interface can be provided to define a service-monitoring dashboard based on user input allowing different users to each create a customized service-monitoring dashboard. Users can select an image for the service-monitoring dashboard (e.g., image for the background of a service-monitoring dashboard, image for an entity and/or service for service-monitoring dashboard), draw a flow chart or a representation of an environment (e.g., IT environment), specify which KPIs to include in the service-monitoring dashboard, configure a KPI widget for each specified KPI, and add one or more adhoc KPI searches to the service-monitoring dashboard. Implementations of the present disclosure provide users with service monitoring information that can be continuously and/or periodically updated. Each service-monitoring dashboard can provide a service-level perspective of how one or more services are performing to help users make operating decisions and/or further evaluate the performance of one or more services.

Implementations are described for a visual interface that displays time-based graphical visualizations that each corresponds to a different KPI reflecting how a service provided by one or more entities is performing. This visual interface may be referred to as a "deep dive." As described herein, machine data pertaining to one or more entities that provide a given service can be presented and viewed in a number of ways. The deep dive visual interface allows an in-depth look at KPI data that reflects how a service or entity is performing over a certain period of time. By having multiple graphical visualizations, each representing a different service or a different aspect of the same service, the deep dive visual interface allows a user to visually correlate the respective KPIs over a defined period of time. In one implementation, the graphical visualizations are all calibrated to the same time scale, so that the values of different KPIs can be compared at any given point in time. In one implementation, the graphical visualizations are all calibrated to different time scales. Although each graphical visualization is displayed in the same visual interface, one or more of the graphical visualizations may have a different time scale than the other graphical visualizations. The different time scale may be more appropriate for the underlying KPI data associated with the one or more graphical visualizations. In one implementation, the graphical visualizations are displayed in parallel lanes, which simplifies visual correlation and allows a user to relate the performance of one service or one aspect of the service (as represented by the KPI values) to the performance of one or more additional services or one or more additional aspects of the same service.

FIG. 1 illustrates a block diagram of an example service provided by entities, in accordance with one or more implementations of the present disclosure. One or more entities **104A, 104B** provide service **102**. An entity **104A, 104B** can be a component in an IT environment. Examples of an entity can include, and are not limited to a host machine, a virtual machine, a switch, a firewall, a router, a sensor, etc. For example, the service **102** may be a web hosting service, and the entities **104A, 104B** may be web servers running on one or more host machines to provide the web hosting service. In another example, an entity could represent a single process on different (physical or virtual) machines. In another example, an entity could represent communication between two different machines.

The service **102** can be monitored using one or more KPIs **106** for the service. A KPI is a type of performance measurement. One or more KPIs can be defined for a service. In the illustrated example, three KPIs **106A-C** are defined for service **102**. KPI **106A** may be a measurement of CPU (central processing unit) usage for the service **102**. KPI **106B** may be a measurement of memory usage for the service **102**. KPI **106C** may be a measurement of request response time for the service **102**.

In one implementation, KPI **106A-C** is derived based on machine data pertaining to entities **104A** and **104B** that provide the service **102** that is associated with the KPI **106A-C**. In another implementation, KPI **106A-C** is derived based on machine data pertaining to entities other than and/or in addition to entities **104A** and **104B**. In another implementation, input (e.g., user input) may be received that defines a custom query, which does not use entity filtering, and is treated as a KPI. Machine data pertaining to a specific entity can be machine data produced by that entity or machine data about that entity, which is produced by another entity. For example, machine data pertaining to entity **104A** can be derived from different sources that may be hosted by entity **104A** and/or some other entity or entities.

A source of machine data can include, for example, a software application, a module, an operating system, a script, an application programming interface, etc. For example, machine data **110B** may be log data that is produced by the operating system of entity **104A**. In another example, machine data **110C** may be produced by a script that is executing on entity **104A**. In yet another example, machine data **110A** may be about an entity **104A** and produced by a software application **120A** that is hosted by another entity to monitor the performance of the entity **104A** through an application programming interface (API).

For example, entity **104A** may be a virtual machine and software application **120A** may be executing outside of the virtual machine (e.g., on a hypervisor or a host operating system) to monitor the performance of the virtual machine via an API. The API can generate network packet data including performance measurements for the virtual machine, such as, memory utilization, CPU usage, etc.

Similarly, machine data pertaining to entity **104B** may include, for example, machine data **110D**, such as log data produced by the operating system of entity **104B**, and machine data **110E**, such as network packets including http responses generated by a web server hosted by entity **104B**.

Implementations of the present disclosure provide for an association between an entity (e.g., a physical machine) and machine data pertaining to that entity (e.g., machine data produced by different sources hosted by the entity or machine data about the entity that may be produced by sources hosted by some other entity or entities). The association may be provided via an entity definition that identifies machine data

from different sources and links the identified machine data with the actual entity to which the machine data pertains, as will be discussed in more detail below in conjunction with FIG. 3 and FIGS. 6-10. Entities that are part of a particular service can be further grouped via a service definition that specifies entity definitions of the entities providing the service, as will be discussed in more detail below in conjunction with FIGS. 11-31.

In the illustrated example, an entity definition for entity **104A** can associate machine data **110A, 110B** and **110C** with entity **104A**, an entity definition for entity **104B** can associate machine data **110D** and **110E** with entity **104B**, and a service definition for service **102** can group entities **104A** and **104B** together, thereby defining a pool of machine data that can be operated on to produce KPIs **106A, 106B** and **106C** for the service **102**. In particular, each KPI **106A, 106B, 106C** of the service **102** can be defined by a search query that produces a value **108A, 108B, 108C** derived from the machine data **110A-E**. As will be discussed in more detail below, according to one implementation, the machine data **110A-E** is identified in entity definitions of entities **104A** and **104B**, and the entity definitions are specified in a service definition of service **102** for which values **108A-C** are produced to indicate how the service **102** is performing at a point in time or during a period of time. For example, KPI **106A** can be defined by a search query that produces value **108A** indicating how the service **102** is performing with respect to CPU usage. KPI **106B** can be defined by a different search query that produces value **108B** indicating how the service **102** is performing with respect to memory usage. KPI **106C** can be defined by yet another search query that produces value **108C** indicating how the service **102** is performing with respect to request response time.

The values **108A-C** for the KPIs can be produced by executing the search query of the respective KPI. In one example, the search query defining a KPI **106A-C** can be executed upon receiving a request (e.g., user request). For example, a service-monitoring dashboard, which is described in greater detail below in conjunction with FIG. 35, can display KPI widgets providing a numerical or graphical representation of the value **108** for a respective KPI **106**. A user may request the service-monitoring dashboard to be displayed at a point in time, and the search queries for the KPIs **106** can be executed in response to the request to produce the value **108** for the respective KPI **106**. The produced values **108** can be displayed in the service-monitoring dashboard.

In another example, the search query defining a KPI **106A-C** can be executed in real-time (continuous execution until interrupted). For example, a user may request the service-monitoring dashboard to be displayed, and the search queries for the KPIs **106** can be executed in response to the request to produce the value **108** for the respective KPI **106**. The produced values **108** can be displayed in the service-monitoring dashboard. The search queries for the KPIs **106** can be continuously executed until interrupted and the values for the search queries can be refreshed in the service-monitoring dashboard with each execution. Examples of interruption can include changing graphical interfaces, stopping execution of a program, etc.

In another example, the search query defining a KPI **106** can be executed based on a schedule. For example, the search query for a KPI (e.g., KPI **106A**) can be executed at one or more particular times (e.g., 6:00 am, 12:00 pm, 6:00 pm, etc.) and/or based on a period of time (e.g., every 5 minutes). In one example, the values (e.g., values **108A**) produced by a search query for a KPI (e.g., KPI **106A**) by executing the search query on a schedule are stored in a data store, and are used to

11

calculate an aggregate KPI score for a service (e.g., service **102**), as described in greater detail below in conjunction with FIGS. **32-33**. An aggregate KPI score for the service **102** is indicative of an overall performance of the KPIs **106** of the service.

In one implementation, the machine data (e.g., machine data **110A-E**) used by a search query defining a KPI (e.g., KPI **106A**) to produce a value can be based on a time range. The time range can be a user-defined time range or a default time range. For example, in the service-monitoring dashboard example above, a user can select, via the service-monitoring dashboard, a time range to use to further specify, for example, based on time-stamps, which machine data should be used by a search query defining a KPI. For example, the time range can be defined as “Last 15 minutes,” which would represent an aggregation period for producing the value. In other words, if the query is executed periodically (e.g., every 5 minutes), the value resulting from each execution can be based on the last 15 minutes on a rolling basis, and the value resulting from each execution can be, for example, the maximum value during a corresponding 15-minute time range, the minimum value during the corresponding 15-minute time range, an average value for the corresponding 15-minute time range, etc.

In another implementation, the time range is a selected (e.g., user-selected) point in time and the definition of an individual KPI can specify the aggregation period for the respective KPI. By including the aggregation period for an individual KPI as part of the definition of the respective KPI, multiple KPIs can run on different aggregation periods, which can more accurately represent certain types of aggregations, such as, distinct counts and sums, improving the utility of defined thresholds. In this manner, the value of each KPI can be displayed at a given point in time. In one example, a user may also select “real time” as the point in time to produce the most up to date value for each KPI using its respective individually defined aggregation period.

An event-processing system can process a search query that defines a KPI of a service. An event-processing system can aggregate heterogeneous machine-generated data (machine data) received from various sources (e.g., servers, databases, applications, networks, etc.) and optionally provide filtering such that data is only represented where it pertains to the entities providing the service. In one example, a KPI may be defined by a user-defined custom query that does not use entity filtering. The aggregated machine data can be processed and represented as events. An event can be represented by a data structure that is associated with a certain point in time and comprises a portion of raw machine data (i.e., machine data). Events are described in greater detail below in conjunction with FIG. **72**. The event-processing system can be configured to perform real-time indexing of the machine data and to execute real-time, scheduled, or historic searches on the source data. An exemplary event-processing system is described in greater detail below in conjunction with FIG. **71**. Example Service Monitoring System

FIG. **2** is a block diagram **200** of one implementation of a service monitoring system **210** for monitoring performance of one or more services using key performance indicators derived from machine data, in accordance with one or more implementations of the present disclosure. The service monitoring system **210** can be hosted by one or more computing machines and can include components for monitoring performance of one or more services. The components can include, for example, an entity module **220**, a service module **230**, a key performance indicator module **240**, a user interface (UI) module **250**, a dashboard module **260**, a deep dive module

12

**270**, and a home page module **280**. The components can be combined together or separated in further components, according to a particular embodiment. The components and/or combinations of components can be hosted on a single computing machine and/or multiple computing machines. The components and/or combinations of components can be hosted on one or more client computing machines and/or server computing machines.

The entity module **220** can create entity definitions. “Create” hereinafter includes “edit” throughout this document. An entity definition is a data structure that associates an entity (e.g., entity **104A** in FIG. **1**) with machine data (e.g., machine data **110A-C** in FIG. **1**). The entity module **220** can determine associations between machine data and entities, and can create an entity definition that associates an individual entity with machine data produced by different sources hosted by that entity and/or other entity(ies). In one implementation, the entity module **220** automatically identifies the entities in an environment (e.g., IT environment), automatically determines, for each entity, which machine data is associated with that particular entity, and automatically generates an entity definition for each entity. In another implementation, the entity module **220** receives input (e.g., user input) for creating an entity definition for an entity, as will be discussed in greater detail below in conjunction with FIGS. **5-10**.

FIG. **3** is a block diagram **300** illustrating an entity definition for an entity, in accordance with one or more implementations of the present disclosure. The entity module **220** can create entity definition **350** that associates an entity **304** with machine data (e.g., machine data **310A**, machine data **310B**, machine data **310C**) pertaining to that entity **304**. Machine data that pertains to a particular entity can be produced by different sources **315** and may be produced in different data formats **330**. For example, the entity **304** may be a host machine that is executing a server application **334** that produces machine data **310B** (e.g., log data). The entity **304** may also host a script **336**, which when executed, produces machine data **310C**. A software application **330**, which is hosted by a different entity (not shown), can monitor the entity **304** and use an API **333** to produce machine data **310A** about the entity **304**.

Each of the machine data **310A-C** can include an alias that references the entity **304**. At least some of the aliases for the particular entity **304** may be different from each other. For example, the alias for entity **304** in machine data **310A** may be an identifier (ID) number **315**, the alias for entity **304** in machine data **310B** may be a hostname **317**, and the alias for entity **304** in machine data **310C** may be an IP (internet protocol) address **319**.

The entity module **220** can receive input for an identifying name **360** for the entity **304** and can include the identifying name **360** in the entity definition **350**. The identifying name **360** can be defined from input (e.g., user input). For example, the entity **304** may be a web server and the entity module **220** may receive input specifying `webserver01.splunk.com` as the identifying name **360**. The identifying name **360** can be used to normalize the different aliases of the entity **304** from the machine data **310A-C** to a single identifier.

A KPI, for example, for monitoring CPU usage for a service provided by the entity **304**, can be defined by a search query directed to search machine data **310A-C** based a service definition, which is described in greater detail below in conjunction with FIG. **4**, associating the entity definition **350** with the KPI, the entity definition **350** associating the entity **304** with the identifying name **360**, and associating the iden-

tifying name **360** (e.g., `webserver01.splunk.com`) with the various aliases (e.g., ID number **315**, hostname **317**, and IP address **319**).

Referring to FIG. 2, the service module **230** can create service definitions for services. A service definition is a data structure that associates one or more entities with a service. The service module **230** can receive input (e.g., user input) of a title and/or description for a service definition. FIG. 4 is a block diagram illustrating a service definition that associates one or more entities with a service, in accordance with one or more implementations of the present disclosure. In another implementation, a service definition specifies one or more other services which a service depends upon and does not associate any entities with the service, as described in greater detail below in conjunction with FIG. 18. In another implementation, a service definition specifies a service as a collection of one or more other services and one or more entities.

In one example, a service **402** is provided by one or more entities **404A-N**. For example, entities **404A-N** may be web servers that provide the service **402** (e.g., web hosting service). In another example, a service **402** may be a database service that provides database data to other services (e.g., analytical services). The entities **404A-N**, which provides the database service, may be database servers.

The service module **230** can include an entity definition **450A-450N**, for a corresponding entity **404A-N** that provides the service **402**, in the service definition **460** for the service **402**. The service module **230** can receive input (e.g., user input) identifying one or more entity definitions to include in a service definition.

The service module **230** can include dependencies **470** in the service definition **460**. The dependencies **470** indicate one or more other services for which the service **402** is dependent upon. For example, another set of entities (e.g., host machines) may define a testing environment that provides a sandbox service for isolating and testing untested programming code changes. In another example, a specific set of entities (e.g., host machines) may define a revision control system that provides a revision control service to a development organization. In yet another example, a set of entities (e.g., switches, firewall systems, and routers) may define a network that provides a networking service. The sandbox service can depend on the revision control service and the networking service. The revision control service can depend on the networking service. If the service **402** is the sandbox service and the service definition **460** is for the sandbox service **402**, the dependencies **470** can include the revision control service and the networking service. The service module **230** can receive input specifying the other service(s) for which the service **402** is dependent on and can include the dependencies **470** between the services in the service definition **460**. In one implementation, the service associated defined by the service definition **460** may be designated as a dependency for another service, and the service definition **460** can include information indicating the other services which depend on the service described by the service definition **460**.

Referring to FIG. 2, the KPI module **240** can create one or more KPIs for a service and include the KPIs in the service definition. For example, in FIG. 4, various aspects (e.g., CPU usage, memory usage, response time, etc.) of the service **402** can be monitored using respective KPIs. The KPI module **240** can receive input (e.g., user input) defining a KPI for each aspect of the service **402** to be monitored and include the KPIs (e.g., KPIs **406A-406N**) in the service definition **460** for the service **402**. Each KPI can be defined by a search query that can produce a value. For example, the KPI **406A** can be

defined by a search query that produces value **408A**, and the KPI **406N** can be defined by a search query that produces value **408N**.

The KPI module **240** can receive input specifying the search processing language for the search query defining the KPI. The input can include a search string defining the search query and/or selection of a data model to define the search query. Data models are described in greater detail below in conjunction with FIGS. 74B-D. The search query can produce, for a corresponding KPI, value **408A-N** derived from machine data that is identified in the entity definitions **450A-N** that are identified in the service definition **460**.

The KPI module **240** can receive input to define one or more thresholds for one or more KPIs. For example, the KPI module **240** can receive input defining one or more thresholds **410A** for KPI **406A** and input defining one or more thresholds **410N** for KPI **406N**. Each threshold defines an end of a range of values representing a certain state for the KPI. Multiple states can be defined for the KPI (e.g., unknown state, trivial state, informational state, normal state, warning state, error state, and critical state), and the current state of the KPI depends on which range the value, which is produced by the search query defining the KPI, falls into. The KPI module **240** can include the threshold definition(s) in the KPI definitions. The service module **230** can include the defined KPIs in the service definition for the service.

The KPI module **240** can calculate an aggregate KPI score **480** for the service for continuous monitoring of the service. The score **480** can be a calculated value **482** for the aggregate of the KPIs for the service to indicate an overall performance of the service. For example, if the service has 10 KPIs and if the values produced by the search queries for 9 of the 10 KPIs indicate that the corresponding KPI is in a normal state, then the value **482** for an aggregate KPI may indicate that the overall performance of the service is satisfactory. Some implementations of calculating a value for an aggregate KPI for the service are discussed in greater detail below in conjunction with FIGS. 32-33.

Referring to FIG. 2, the service monitoring system **210** can be coupled to one or more data stores **290**. The entity definitions, the service definitions, and the KPI definitions can be stored in the data store(s) **290** that are coupled to the service monitoring system **210**. The entity definitions, the service definitions, and the KPI definitions can be stored in a data store **290** in a key-value store, a configuration file, a lookup file, a database, or in metadata fields associated with events representing the machine data. A data store **290** can be a persistent storage that is capable of storing data. A persistent storage can be a local storage unit or a remote storage unit. Persistent storage can be a magnetic storage unit, optical storage unit, solid state storage unit, electronic storage units (main memory), or similar storage unit. Persistent storage can be a monolithic device or a distributed set of devices. A 'set', as used herein, refers to any positive whole number of items.

The user interface (UI) module **250** can generate graphical interfaces for creating and/or editing entity definitions for entities, creating and/or editing service definitions for services, defining key performance indicators (KPIs) for services, setting thresholds for the KPIs, and defining aggregate KPI scores for services. The graphical interfaces can be user interfaces and/or graphical user interfaces (GUIs).

The UI module **250** can cause the display of the graphical interfaces and can receive input via the graphical interfaces. The entity module **220**, service module **230**, KPI module **240**, dashboard module **260**, deep dive module **270**, and home page module **280** can receive input via the graphical interfaces generated by the UI module **250**. The entity module



15

220, service module 230, KPI module 240, dashboard module 260, deep dive module 270, and home page module 280 can provide data to be displayed in the graphical interfaces to the UI module 250, and the UI module 250 can cause the display of the data in the graphical interfaces.

The dashboard module 260 can create a service-monitoring dashboard. In one implementation, dashboard module 260 works in connection with UI module 250 to present a dashboard-creation graphical interface that includes a modifiable dashboard template, an interface containing drawing tools to customize a service-monitoring dashboard to define flow charts, text and connections between different elements on the service-monitoring dashboard, a KPI-selection interface and/or service selection interface, and a configuration interface for creating service-monitoring dashboard. The service-monitoring dashboard displays one or more KPI widgets. Each KPI widget can provide a numerical or graphical representation of one or more values for a corresponding KPI indicating how an aspect of a service is performing at one or more points in time. Dashboard module 260 can work in connection with UI module 250 to define the service-monitoring dashboard in response to user input, and to cause display of the service-monitoring dashboard including the one or more KPI widgets. The input can be used to customize the service-monitoring dashboard. The input can include for example, selection of one or more images for the service-monitoring dashboard (e.g., a background image for the service-monitoring dashboard, an image to represent an entity and/or service), creation and representation of adhoc search in the form of KPI widgets, selection of one or more KPIs to represent in the service-monitoring dashboard, selection of a KPI widget for each selected KPI. The input can be stored in the one or more data stores 290 that are coupled to the dashboard module 260. In other implementations, some other software or hardware module may perform the actions associated with generating and displaying the service-monitoring dashboard, although the general functionality and features of the service-monitoring dashboard should remain as described herein. Some implementations of creating the service-monitoring dashboard and causing display of the service-monitoring dashboard are discussed in greater detail below in conjunction with FIGS. 35-47.

In one implementation, deep dive module 270 works in connection with UI module 250 to present a wizard for creation and editing of the deep dive visual interface, to generate the deep dive visual interface in response to user input, and to cause display of the deep dive visual interface including the one or more graphical visualizations. The input can be stored in the one or more data stores 290 that are coupled to the deep dive module 270. In other implementations, some other software or hardware module may perform the actions associated with generating and displaying the deep dive visual interface, although the general functionality and features of deep dive should remain as described herein. Some implementations of creating the deep dive visual interface and causing display of the deep dive visual interface are discussed in greater detail below in conjunction with FIGS. 49-70.

The home page module 280 can create a home page graphical interface. The home page graphical interface can include one or more tiles, where each tile represents a service-related alarm, service-monitoring dashboard, a deep dive visual interface, or the value of a particular KPI. In one implementation home page module 280 works in connection with UI module 250. The UI module 250 can cause the display of the home page graphical interface. The home page module 280 can receive input (e.g., user input) to request a service-monitoring dashboard or a deep dive to be displayed. The input can

16

include for example, selection of a tile representing a service-monitoring dashboard or a deep dive. In other implementations, some other software or hardware module may perform the actions associated with generating and displaying the home page graphical interface, although the general functionality and features of the home page graphical interface should remain as described herein. An example home page graphical interface is discussed in greater detail below in conjunction with FIG. 48.

Referring to FIG. 2, the service monitoring system 210 can be coupled to an event processing system 205 via one or more networks. The event processing system 205 can receive a request from the service monitoring system 210 to process a search query. For example, the dashboard module 260 may receive input request to display a service-monitoring dashboard with one or more KPI widgets. The dashboard module 260 can request the event processing system 205 to process a search query for each KPI represented by a KPI widget in the service-monitoring dashboard. Some implementations of an event processing system 205 are discussed in greater detail below in conjunction with FIG. 71.

The one or more networks can include one or more public networks (e.g., the Internet), one or more private networks (e.g., a local area network (LAN) or one or more wide area networks (WAN)), one or more wired networks (e.g., Ethernet network), one or more wireless networks (e.g., an 802.11 network or a Wi-Fi network), one or more cellular networks (e.g., a Long Term Evolution (LTE) network), routers, hubs, switches, server computers, and/or a combination thereof.

Key Performance Indicators

FIG. 5 is a flow diagram of an implementation of a method 500 for creating one or more key performance indicators for a service, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, at least a portion of method is performed by a client computing machine. In another implementation, at least a portion of method is performed by a server computing machine.

At block 502, the computing machine creates one or more entity definitions, each for a corresponding entity. Each entity definition associates an entity with machine data that pertains to that entity. As described above, various machine data may be associated with a particular entity, but may use different aliases for identifying the same entity. The entity definition for an entity normalizes the different aliases of that entity. In one implementation, the computing machine receives input for creating the entity definition. The input can be user input. Some implementations of creating an entity definition for an entity from input received via a graphical user interface are discussed in greater detail below in conjunction with FIGS. 6-10.

In another implementation, the computing machine imports a data file (e.g., CSV (comma-separated values) data file) that includes information identifying entities in an environment and uses the data file to automatically create entity definitions for the entities described in the data file. The data file may be stored in a data store (e.g., data store 290 in FIG. 2) that is coupled to the computing machine.

In another implementation, the computing machine automatically (without any user input) identifies one or more aliases for an entity in machine data, and automatically creates an entity definition in response to automatically identifying the aliases of the entity in the machine data. For example, the computing machine can execute a search query

17

from a saved search to extract data to identify an alias for an entity in machine data from one or more sources, and automatically create an entity definition for the entity based on the identified aliases. Some implementations of creating an entity definition from importing a data file and/or from a saved search are discussed in greater detail below in conjunction with FIG. 16.

At block 504, the computing machine creates a service definition for a service using the entity definitions of the one or more entities that provide the service, according to one implementation. A service definition can relate one or more entities to a service. For example, the service definition can include an entity definition for each of the entities that provide the service. In one implementation, the computing machine receives input (e.g., user input) for creating the service definition. Some implementations of creating a service definition from input received via a graphical interface are discussed in more detail below in conjunction with FIGS. 11-18. In one implementation, the computing machine automatically creates a service definition for a service. In another example, a service may not directly be provided by one or more entities, and the service definition for the service may not directly relate one or more entities to the service. For example, a service definition for a service may not contain any entity definitions and may contain information indicating that the service is dependent on one or more other services. A service that is dependent on one or more other services is described in greater detail below in conjunction with FIG. 18. For example, a business service may not be directly provided by one or more entities and may be dependent on one or more other services. For example, an online store service may depend on an e-commerce service provided by an e-commerce system, a database service, and a network service. The online store service can be monitored via the entities of the other services (e.g., e-commerce service, database service, and network service) upon which the service depends on.

At block 506, the computing machine creates one or more key performance indicators (KPIs) corresponding to one or more aspects of the service. An aspect of a service may refer to a certain characteristic of the service that can be measured at various points in time during the operation of the service. For example, aspects of a web hosting service may include request response time, CPU usage, and memory usage. Each KPI for the service can be defined by a search query that produces a value derived from the machine data that is identified in the entity definitions included in the service definition for the service. Each value is indicative of how an aspect of the service is performing at a point in time or during a period of time. In one implementation, the computing machine receives input (e.g., user input) for creating the KPI(s) for the service. Some implementations of creating KPI(s) for a service from input received via a graphical interface will be discussed in greater detail below in conjunction with FIGS. 19-31. In one implementation, the computing machine automatically creates one or more key performance indicators (KPIs) corresponding to one or more aspects of the service.

FIG. 6 is a flow diagram of an implementation of a method 600 for creating an entity definition for an entity, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, at least a portion of method is performed by

18

a client computing machine. In another implementation, at least a portion of method is performed by a server computing machine.

At block 602, the computing machine receives input of an identifying name for referencing the entity definition for an entity. The input can be user input. The user input can be received via a graphical interface. Some implementations of creating an entity definition via input received from a graphical interface are discussed in greater detail below in conjunction with FIGS. 7-10. The identifying name can be a unique name.

At block 604, the computing machine receives input (e.g., user input) specifying one or more search fields ("fields") representing the entity in machine data from different sources, to be used to normalize different aliases of the entity. Machine data can be represented as events. As described above, the computing machine can be coupled to an event processing system (e.g., event processing system 205 in FIG. 2). The event processing system can process machine data to represent the machine data as events. Each of the events is raw data, and when a late binding schema is applied to the events, values for fields defined by the schema are extracted from the events. A number of "default fields" that specify metadata about the events rather than data in the events themselves can be created automatically. For example, such default fields can specify: a timestamp for the event data; a host from which the event data originated; a source of the event data; and a source type for the event data. These default fields may be determined automatically when the events are created, indexed or stored. Each event has metadata associated with the respective event. Implementations of the event processing system processing the machine data to be represented as events are discussed in greater detail below in conjunction with FIG. 71.

At block 606, the computing machine receives input (e.g., user input) specifying one or more search values ("values") for the fields to establish associations between the entity and machine data. The values can be used to search for the events that have matching values for the above fields. The entity can be associated with the machine data that is represented by the events that have fields that store values that match the received input.

The computing machine can optionally also receive input (e.g., user input) specifying a type of entity to which the entity definition applies. The computing machine can optionally also receive input (e.g., user input) associating the entity of the entity definition with one or more services. Some implementations of receiving input for an entity type for an entity definition and associating the entity with one or more services are discussed in greater detail below in conjunction with FIGS. 9A-B.

FIG. 7 illustrates an example of a GUI 700 of a service monitoring system for creating and/or editing entity definition(s) and/or service definition(s), in accordance with one or more implementations of the present disclosure. One or more GUIs of the service monitoring system can include GUI elements to receive input and to display data. The GUI elements can include, for example, and are not limited to, a text box, a button, a link, a selection button, a drop down menu, a sliding bar, a selection button, an input field, etc. In one implementation, GUI 700 includes a menu item, such as Configure 702, to facilitate the creation of entity definitions and service definitions.

Upon the selection of the Configure 702 menu item, a drop-down menu 704 listing configuration options can be displayed. If the user selects the entities option 706 from the drop-down menu 704, a GUI for creating an entity definition can be displayed, as discussed in more detail below in con-

19

junction with FIG. 8. If the user selects the services option **708** from the drop-down menu **704**, a GUI for creating a service definition can be displayed, as discussed in more detail below in conjunction with FIG. 11.

FIG. 8 illustrates an example of a GUI **800** of a service monitoring system for creating and/or editing entity definitions, in accordance with one or more implementations of the present disclosure. GUI **800** can display a list **802** of entity definitions that have already been created. Each entity definition in the list **802** can include a button **804** for requesting a drop-down menu **810** listing editing options to edit the corresponding entity definition. Editing can include editing the entity definition and/or deleting the entity definition. When an editing option is selected from the drop-down menu **810**, one or more additional GUIs can be displayed for editing the entity definition. GUI **800** can include an import button **806** for importing a data file (e.g., CSV file) for auto-discovery of entities and automatic generation of entity definitions for the discovered entities. The data file can include a list of entities that exist in an environment (e.g., IT environment). The service monitoring system can use the data file to automatically create an entity definition for an entity in the list. In one implementation, the service monitoring system uses the data file to automatically create an entity definition for each entity in the list. GUI **800** can include a button **808** that a user can activate to proceed to the creation of an entity definition, which leads to GUI **900** of FIG. 9A. The automatic generation of entity definitions for entities is described in greater detail below in conjunction with FIG. 16.

FIG. 9A illustrates an example of a GUI **900** of a service monitoring system for creating an entity definition, in accordance with one or more implementations of the present disclosure. GUI **900** can facilitate user input specifying an identifying name **904** for the entity, an entity type **906** for the entity, field(s) **908** and value(s) **910** for the fields **908** to use during the search to find events pertaining to the entity, and any services **912** that the entity provides. The entity type **906** can describe the particular entity. For example, the entity may be a host machine that is executing a webserver application that produces machine data. FIG. 9B illustrates an example of input received via GUI **900** for creating an entity definition, in accordance with one or more implementations of the present disclosure.

For example, the identifying name **904** is `webserver01.splunk.com` and the entity type **906** is web server. Examples of entity type can include, and are not limited to, host machine, virtual machine, type of server (e.g., web server, email server, database server, etc.) switch, firewall, router, sensor, etc. The fields **908** that are part of the entity definition can be used to normalize the various aliases for the entity. For example, the entity definition specifies three fields **920, 922, 924** and four values **910** (e.g., values **930, 932, 934, 936**) to associate the entity with the events that include any of the four values in any of the three fields.

For example, the event processing system (e.g., event processing system **205** in FIG. 2) can apply a late-binding schema to the events to extract values for fields (e.g., host field, ip field, and dest field) defined by the schema and determine which events have values that are extracted for a host field that includes `10.11.12.13`, `webserver01.splunk.com`, `webserver01`, or `vm-0123`, determine which events have values that are extracted for an ip field that includes `10.11.12.13`, `webserver01.splunk.com`, `webserver01`, or `vm-0123`, or a dest field that includes `10.11.12.13`, `webserver01.splunk.com`, `webserver01`, or `vm-0123`. The machine data that relates to the events that are

20

produced from the search is the machine data that is associated with the entity `webserver01.splunk.com`.

In another implementation, the entity definition can specify one or more values **910** to use for a specific field **908**. For example, the value **930** (`10.11.12.13`) may be used for extracting values for the ip field and determine which values match the value **930**, and the value **932** (`webserver01.splunk.com`) and the value **936** (`vm-0123`) may be used for extracting values for the host **920** field and determining which values match the value **932** or value **936**.

In another implementation, GUI **900** includes a list of identifying field/value pairs. A search term that is modeled after these entities can be constructed, such that, when a late-binding schema is applied to events, values that match the identifiers associated with the fields defined by the schema will be extracted. For example, if `identifier.fields="X,Y"` then the entity definition should include input specifying fields labeled "X" and "Y". The entity definition should also include input mapping the fields. For example, the entity definition can include the mapping of the fields as `"X":"1", "Y":["2","3"]`. The event processing system (e.g., event processing system **205** in FIG. 2) can apply a late-binding schema to the events to extract values for fields (e.g., X and Y) defined by the schema and determine which events have values extracted for an X field that include "1", or which events have values extracted for a Y field that include "2", or which events have values extracted for a Y field that include "3".

GUI **900** can facilitate user input specifying any services **912** that the entity provides. The input can specify one or more services that have corresponding service definitions. For example, if there is a service definition for a service named web hosting service that is provided by the entity corresponding to the entity definition, then a user can specify the web hosting service as a service **912** in the entity definition.

The save button **916** can be selected to save the entity definition in a data store (e.g., data store **290** in FIG. 2). The saved entity definition can be edited. FIG. 10 illustrates an example of a GUI **1000** of a service monitoring system for creating and/or editing entity definitions, in accordance with one or more implementations of the present disclosure. GUI **1000** can display a list **1002** of entity definitions that have already been created. For example, list **1002** includes the entity definition `webserver01.splunk.com` that can be selected for editing.

FIG. 11 is a flow diagram of an implementation of a method **1100** for creating a service definition for a service, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, at least a portion of method is performed by a client computing machine. In another implementation, at least a portion of method is performed by a server computing machine.

At block **1102**, the computing machine receives input of a title for referencing a service definition for a service. At block **1104**, the computing machine receives input identifying one or more entities providing the service and associates the identified entities with the service definition of the service at block **1106**.

At block **1108**, the computing machine creates one or more key performance indicators for the service and associates the key performance indicators with the service definition of the service at block **1110**. Some implementations of creating one

## 21

or more key performance indicators are discussed in greater detail below in conjunction with FIGS. 19-31.

At block 1112, the computing machine receives input identifying one or more other services which the service is dependent upon and associates the identified other services with the service definition of the service at block 1114. The computing machine can include an indication in the service definition that the service is dependent on another service for which a service definition has been created.

At block 1116, the computing machine can optionally define an aggregate KPI score to be calculated for the service to indicate an overall performance of the service. The score can be a value for an aggregate of the KPIs for the service. The aggregate KPI score can be periodically calculated for continuous monitoring of the service. For example, the aggregate KPI score for a service can be updated in real-time (continuously updated until interrupted). In one implementation, the aggregate KPI score for a service is updated periodically (e.g., every second). Some implementations of determining an aggregate KPI score for the service are discussed in greater detail below in conjunction with FIGS. 32-34.

FIG. 12 illustrates an example of a GUI 1200 of a service monitoring system for creating and/or editing service definitions, in accordance with one or more implementations of the present disclosure. GUI 1200 can display a list 1202 of service definitions that have already been created. Each service definition in the list 1202 can include a button 1204 to proceed to a drop-down menu 1208 listing editing options related to the corresponding service definition. Editing options can include editing the service definition, editing one or more KPIs for the service, editing a title and/or description of the service description, and/or deleting the service definition. When an editing option is selected from the drop-down menu 1208, one or more other GUIs can be displayed for editing the service definition. GUI 1200 can include a button 1210 to proceed to the creation of a new service definition.

FIG. 13 illustrates an example of a GUI 1300 of a service monitoring system for creating a service definition, in accordance with one or more implementations of the present disclosure. GUI 1300 can facilitate user input specifying a title 1302 and optionally a description 1304 for the service definition for a service. GUI 1300 can include a button 1306 to proceed to GUI 1400 of FIG. 14, for associating entities with the service, creating KPIs for the service, and indicating dependencies for the service.

FIG. 14 illustrates an example of a GUI 1400 of a service monitoring system for defining elements of a service definition, in accordance with one or more implementations of the present disclosure. GUI 1400 can include an accordion pane (accordion section) 1402, which when selected, displays fields for facilitating input for creating and/or editing a title 1404 of a service definition, and input for a description 1406 of the service that corresponds to the service definition. If input for the title 1404 and/or description 1406 was previously received, for example, from GUI 1300 in FIG. 13, GUI 1400 can display the title 1404 and description 1406.

GUI 1400 can include a drop-down 1410 for receiving input for creating one or more KPIs for the service. If the drop-down 1410 is selected, GUI 1900 in FIG. 19 is displayed as described in greater detail below.

GUI 1400 can include a drop-down 1412 for receiving input for specifying dependencies for the service. If the drop-down 1412 is selected, GUI 1800 in FIG. 18 is displayed as described in greater detail below.

GUI 1400 can include one or more buttons 1408 to specify whether entities are associated with the service. A selection of "No" 1416 indicates that the service is not associated with any

## 22

entities and the service definition is not associated with any entity definitions. For example, a service may not be associated with any entities if an end user intends to use the service and corresponding service definition for testing purposes and/or experimental purposes. In another example, a service may not be associated with any entities if the service is dependent one or more other services, and the service is being monitored via the entities of the one or more other services upon which the service depends upon. For example, an end user may wish to use a service without entities as a way to track a business service based on the services which the business service depends upon. If "Yes" 1414 is selected, GUI 1500 in FIG. 15 is displayed as described in greater detail below.

FIG. 15 illustrates an example of a GUI 1500 of a service monitoring system for associating one or more entities with a service by associating one or more entity definitions with a service definition, in accordance with one or more implementations of the present disclosure. GUI 1500 can include a button 1510 for creating a new entity definition. If button 1510 is selected, GUI 1600 in FIG. 16 is displayed facilitating user input for creating an entity definition.

FIG. 16 illustrates an example of a GUI 1600 facilitating user input for creating an entity definition, in accordance with one or more implementations of the present disclosure. For example, GUI 1600 can include multiple fields 1601 for creating an entity definition, as discussed above in conjunction with FIG. 6. GUI 1600 can include a button 1603, which when selected can display one or more UIs (e.g., GUIs or command line interface) for importing a data file for creating an entity definition. The data file can be a CSV (comma-separated values) data file that includes information identifying entities in an environment. The data file can be used to automatically create entity definitions for the entities described in the data file. GUI 1600 can include a button 1605, which when selected can display one or more UIs (e.g., GUIs or command line interface) for using a saved search for creating an entity definition. For example, the computing machine can execute a search query from a saved search to extract data to identify an alias for an entity in machine data from one or more sources, and automatically create an entity definition for the entity based on the identified aliases.

Referring to FIG. 15, GUI 1500 can include an availability list 1504 of entity definitions for entities, which can be selected to be associated with the service definition. The availability list 1504 can include one or more entity definitions. For example, the availability list 1504 may include thousands of entity definitions. GUI 1500 can include a filter box 1502 to receive input for filtering the availability list 1504 of entity definitions to display a portion of the entity definitions. Each entity definition in the availability list 1502 can include the entity definition name 1506 and the entity type 1508. GUI 1500 can facilitate user input for selecting an entity definition from the availability list 1504 and dragging the selected entity definition to a selected list 1512 to indicate that the entity for the selected entity definition is associated with service of the service definition. For example, entity definition 1514 (e.g., webserver01.splunk.com) can be selected and dragged to the selected list 1512.

FIG. 17 illustrates an example of a GUI 1700 indicating one or more entities associated with a service based on input, in accordance with one or more implementations of the present disclosure. The selected list 1712 can include the entity definition (e.g., webserver01.splunk.com) that was dragged from the availability list 1704. The availability list 1704 can remove any selected entity definitions (e.g., webserver01.splunk.com). The selected list 1712 indicates

which entities are members of a service via the entity definitions of the entities and service definition for the service.

FIG. 18 illustrates an example of a GUI 1800 of a service monitoring system for specifying dependencies for the service, in accordance with one or more implementations of the present disclosure. GUI 1800 can include an availability list 1804 of services that each has a corresponding service definition. The availability list 1804 can include one or more services. For example, the availability list 1804 may include dozens of services. GUI 1800 can include a filter box 1802 to receive input for filtering the availability list 1804 of services to display a portion of the services. GUI 1800 can facilitate user input for selecting a service from the availability list 1804 and dragging the selected service to a dependent services list 1812 to indicate that the service is dependent on the services in the dependent services list 1812. For example, the service definition may be for a Sandbox service. For example, the drop-down 1801 can be selected to display a title “Sandbox” in the service information for the service definition. The availability list 1804 may initially include four other services: (1) Revision Control service, (2) Networking service, (3) Web Hosting service, and (4) Database service. The Sandbox service may depend on the Revision Control service and the Networking service. A user may select the Revision Control service and Networking service from the availability list 1804 and drag the Revision Control service and Networking service to the dependent services list 1812 to indicate that the Sandbox service is dependent on the Revision Control service and Networking service. In one implementation, GUI 1800 further displays a list of other services which depend on the service described by the service definition that is being created and/or edited.

Thresholds for Key Performance Indicators

FIG. 19 is a flow diagram of an implementation of a method 1900 for creating one or more key performance indicators for a service, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block 1902, the computing machine receives input (e.g., user input) of a name for a KPI to monitor a service or an aspect of the service. For example, a user may wish to monitor the service’s response time for requests, and the name of the KPI may be “Request Response Time.” In another example, a user may wish to monitor the load of CPU(s) for the service, and the name of the KPI may be “CPU Usage.”

At block 1904, the computing machine creates a search query to produce a value indicative of how the service or the aspect of the service is performing. For example, the value can indicate how the aspect (e.g., CPU usage, memory usage, request response time) is performing at point in time or during a period of time. Some implementations for creating a search query are discussed in greater detail below in conjunction with FIG. 20. In one implementation, the computing machine receives input (e.g., user input), via a graphical interface, of search processing language defining the search query. Some implementations for creating a search query from input of search processing language are discussed in greater detail below in conjunction with FIGS. 22-23. In one implementation, the computing machine receives input (e.g., user input) for defining the search query using a data model. Some imple-

mentations for creating a search query using a data model are discussed in greater detail below in conjunction with FIGS. 24-26.

At block 1906, the computing machine sets one or more thresholds for the KPI. Each threshold defines an end of a range of values. Each range of values represents a state for the KPI. The KPI can be in one of the states (e.g., normal state, warning state, critical state) depending on which range the value falls into. Some implementations for setting one or more thresholds for the KPI are discussed in greater detail below in conjunction with FIGS. 28-31.

FIG. 20 is a flow diagram of an implementation of a method 2000 for creating a search query, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block 2002, the computing machine receives input (e.g., user input) specifying a field to use to derive a value indicative of the performance of a service or an aspect of the service to be monitored. As described above, machine data can be represented as events. Each of the events is raw data. A late-binding schema can be applied to each of the events to extract values for fields defined by the schema. The received input can include the name of the field from which to extract a value when executing the search query. For example, the received user input may be the field name “spent” that can be used to produce a value indicating the time spent to respond to a request.

At block 2004, the computing machine optionally receives input specifying a statistical function to calculate a statistic using the value in the field. In one implementation, a statistic is calculated using the value(s) from the field, and the calculated statistic is indicative of how the service or the aspect of the service is performing. As discussed above, the machine data used by a search query for a KPI to produce a value can be based on a time range. For example, the time range can be defined as “Last 15 minutes,” which would represent an aggregation period for producing the value. In other works, if the query is executed periodically (e.g., every 5 minutes), the value resulting from each execution can be based on the last 15 minutes on a rolling basis, and the value resulting from each execution can be based on the statistical function. Examples of statistical functions include, and are not limited to, average, count, count of distinct values, maximum, mean, minimum, sum, etc. For example, the value may be from the field “spent” the time range may be “Last 15 minutes,” and the input may specify a statistical function of average to define the search query that should produce the average of the values of field “spent” for the corresponding 15 minute time range as a statistic. In another example, the value may be a count of events satisfying the search criteria that include a constraint for the field (e.g., if the field is “response time,” and the KPI is focused on measuring the number of slow responses (e.g., “response time” below x) issued by the service).

At block 2006, the computing machine defines the search query based on the specified field and the statistical function. The computing machine may also optionally receive input of an alias to use for a result of the search query. The alias can be used to have the result of the search query to be compared to one or more thresholds assigned to the KPI.

25

FIG. 21 illustrates an example of a GUI 2100 of a service monitoring system for creating a KPI for a service, in accordance with one or more implementations of the present disclosure. GUI 2100 can display a list 2104 of KPIs that have already been created for the service and associated with the service via the service definition. For example, the service definition “Web Hosting” includes a KPI “Storage Capacity” and a KPI “Memory Usage”. GUI 2100 can include a button 2106 for editing a KPI. A KPI in the list 2104 can be selected and the button 2106 can be activated to edit the selected KPI. GUI 2100 can include a button 2102 for creating a new KPI. If button 2102 is activated, GUI 2200 in FIG. 22 is displayed facilitating user input for creating a KPI.

FIG. 22 illustrates an example of a GUI 2200 of a service monitoring system for creating a KPI for a service, in accordance with one or more implementations of the present disclosure. GUI 2200 can facilitate user input specifying a name 2202 and optionally a description 2204 for a KPI for a service. The name 2202 can indicate an aspect of the service that is to be monitored using the KPI. As described above, the KPI is defined by a search query that produces a value derived from machine data pertaining to one or more entities identified in a service definition for the service. The produced value is indicative of how an aspect of the service is performing. In one example, the produced value is the value extracted from a field when the search query is executed. In another example, the produced value is a result from calculating a statistic based on the value in the field.

In one implementation, the search query is defined from input (e.g., user input), received via a graphical interface, of search processing language defining the search query. GUI 2200 can include a button 2206 for facilitating user input of search processing language defining the search query. If button 2206 is selected, a GUI for facilitating user input of search processing language defining the search query can be displayed, as discussed in greater detail below in conjunction with FIG. 23.

Referring to FIG. 22, in another implementation, the search query is defined using a data model. GUI 2200 can include a button 2208 for facilitating user input of a data model for defining the search query. If button 2208 is selected, a GUI for facilitating user input for defining the search query using a data model can be displayed, as discussed in greater detail below in conjunction with FIG. 24.

FIG. 23 illustrates an example of a GUI 2300 of a service monitoring system for receiving input of search processing language for defining a search query for a KPI for a service, in accordance with one or more implementations of the present disclosure. GUI 2300 can facilitate user input specifying a KPI name 2301, which can optionally indicate an aspect of the service to monitor with the KPI, and optionally a description 2302 for a KPI for a service. For example, the aspect of the service to monitor can be response time for received requests, and the KPI name 2301 can be Request Response Time. GUI 2300 can facilitate user input specifying search processing language 2303 that defines the search query for the Request Response Time KPI. The input for the search processing language 2303 can specify a name of a field (e.g., spent 2313) to use to extract a value indicative of the performance of an aspect (e.g., response time) to be monitored for a service. The input of the field (e.g., spent 2313) designates which data to extract from an event when the search query is executed.

The input can optionally specify a statistical function (e.g., avg 2311) that should be used to calculate a statistic based on the value corresponding to a late-binding schema being applied to an event. The late-binding schema will extract a

26

portion of event data corresponding to the field (e.g., spent 2313). For example, the value associated with the field “spent” can be extracted from an event by applying a late-binding schema to the event. The input may specify that the average of the values corresponding to the field “spent” should be produced by the search query. The input can optionally specify an alias (e.g., rsp\_time 2315) to use (e.g., as a virtual field name) for a result of the search query (e.g., avg(spent) 2314). The alias 2315 can be used to have the result of the search query to be compared with one or more thresholds assigned to the KPI.

GUI 2300 can display a link 2304 to facilitate user input to request that the search criteria be tested by running the search query for the KPI. In one implementation, when input is received requesting to test the search criteria for the search query, a search GUI is displayed.

In some implementations, GUI 2300 can facilitate user input for creating one or more thresholds for the KPI. The KPI can be in one of multiple states (e.g., normal, warning, critical). Each state can be represented by a range of values. During a certain time, the KPI can be in one of the states depending on which range the value, which is produced at that time by the search query for the KPI, falls into. GUI 2300 can include a button 2307 for creating the threshold for the KPI. Each threshold for a KPI defines an end of a range of values, which represents one of the states. Some implementations for creating one or more thresholds for the KPI are discussed in greater detail below in conjunction with FIGS. 28-31.

GUI 2300 can include a button 2309 for editing which entity definitions to use for the KPI. Some implementations for editing which entity definitions to use for the KPI are discussed in greater detail below in conjunction with FIG. 27.

In some implementations, GUI 2300 can include a button 2320 to receive input assigning a weight to the KPI to indicate an importance of the KPI for the service relative to other KPIs defined for the service. The weight can be used for calculating an aggregate KPI score for the service to indicate an overall performance for the service, as discussed in greater detail below in conjunction with FIG. 32. GUI 2300 can include a button 2323 to receive input to define how often the KPI should be measured (e.g., how often the search query defining the KPI should be executed) for calculating an aggregate KPI score for the service to indicate an overall performance for the service, as discussed in greater detail below in conjunction with FIG. 32. The importance (e.g., weight) of the KPI and the frequency of monitoring (e.g., a schedule for executing the search query) of the KPI can be used to determine an aggregate KPI score for the service. The score can be a value of an aggregate of the KPIs of the service. Some implementations for using the importance and frequency of monitoring for each KPI to determine an aggregate KPI score for the service are discussed in greater detail below in conjunction with FIGS. 32-33.

GUI 2300 can display an input box 2305 for a field to which the threshold(s) can be applied. In particular, a threshold can be applied to the value produced by the search query defining the KPI. Applying a threshold to the value produced by the search query is described in greater detail below in conjunction with FIG. 29.

FIG. 24 illustrates an example of a GUI 2400 of a service monitoring system for defining a search query for a KPI using a data model, in accordance with one or more implementations of the present disclosure. GUI 2400 can facilitate user input specifying a name 2403 and optionally a description 2404 for a KPI for a service. For example, the aspect of the service to monitor can be CPU utilization, and the KPI name 2403 can be CPU Usage. If button 2402 is selected, GUI 2400

displays button **2406** and button **2408** for defining the search query for the KPI using a data model. A data model refers to one or more objects grouped in a hierarchical manner and can include a root object and, optionally, one or more child objects that can be linked to the root object. A root object can be defined by search criteria for a query to produce a certain set of events, and a set of fields that can be exposed to operate on those events. Each child object can inherit the search criteria of its parent object and can have additional search criteria to further filter out events represented by its parent object. Each child object may also include at least some of the fields of its parent object and optionally additional fields specific to the child object, as will be discussed in greater detail below in conjunction with FIGS. **74B-D**.

If button **2402** is selected, GUI **2500** in FIG. **25** is displayed for facilitating user input for selecting a data model to assist with defining the search query. FIG. **25** illustrates an example of a GUI **2500** of a service monitoring system for facilitating user input for selecting a data model and an object of the data model to use for defining the search query, in accordance with one or more implementations of the present disclosure. GUI **2500** can include a drop-down menu **2503**, which when expanded, displays a list of available data models. When a data model is selected, GUI **2500** can display a list **2505** of objects pertaining to the selected data model. For example, the data model Performance is selected and the objects pertaining to the Performance data model are included in the list **2505**. Objects of a data model are described in greater detail below in conjunction with FIGS. **74B-D**. When an object in the list **2505** is selected, GUI **2500** can display a list **2511** of fields pertaining to the selected object. For example, the CPU object **2509** is selected and the fields pertaining to the CPU object **2509** are included in the list **2511**. GUI **2500** can facilitate user input of a selection of a field in the list **2511**. The selected field (e.g., `cpu_load_percent` **2513**) is the field to use for the search query to derive a value indicative of the performance of an aspect (e.g., CPU usage) of the service. The derived value can be, for example, the field's value extracted from an event when the search query is executed, a statistic calculated based on one or more values of the field in one or more events located when the search query is executed, a count of events satisfying the search criteria that include a constraint for the field (e.g., if the field is "response time" and the KPI is focused on measuring the number of slow responses (e.g., "response time" below x) issued by the service).

Referring to FIG. **24**, GUI **2400** can display a button **2408** for optionally selecting a statistical function to calculate a statistic using the value(s) from the field (e.g., `cpu_load_percent` **2513**). If a statistic is calculated, the result from calculating the statistic becomes the produced value from the search query, which indicates how an aspect of the service is performing. When button **2408** is selected, GUI **2400** can display a drop-down list of statistics. The list of statistics can include, and are not limited to, average, count, count of distinct values, maximum, mean, minimum, sum, etc. For example, a user may select "average" and the value produced by the search query may be the average of the values of field `cpu_load_percent` **2513** for a specified time range (e.g., "Last 15 minutes"). FIG. **26** illustrates an example of a GUI **2600** of a service monitoring system for displaying a selected statistic **2601** (e.g., average), in accordance with one or more implementations of the present disclosure.

Referring to FIG. **24**, GUI **2400** can facilitate user input for creating one or more thresholds for the KPI. GUI **2400** can include a button **2410** for creating the threshold(s) for the KPI. Some implementations for creating one or more thresh-

olds for the KPI are discussed in greater detail below in conjunction with FIGS. **28-31**.

GUI **2400** can include a button **2412** for editing which entity definitions to use for the KPI. Some implementations for editing which entity definitions to use for the KPI are discussed in greater detail below in conjunction with FIG. **27**.

GUI **2400** can include a button **2418** for saving a definition of a KPI and an association of the defined KPI with a service. The KPI definition and association with a service can be stored in a data store.

The value for the KPI can be produced by executing the search query of the KPI. In one example, the search query defining the KPI can be executed upon receiving a request (e.g., user request). For example, a service-monitoring dashboard, which is described in greater detail below in conjunction with FIG. **35**, can display a KPI widget providing a numerical or graphical representation of the value for the KPI. A user may request the service-monitoring dashboard to be displayed, and the computing machine can cause the search query for the KPI to execute in response to the request to produce the value for the KPI. The produced value can be displayed in the service-monitoring dashboard.

In another example, the search query defining the KPI can be executed based on a schedule. For example, the search query for a KPI can be executed at one or more particular times (e.g., 6:00 am, 12:00 pm, 6:00 pm, etc.) and/or based on a period of time (e.g., every 5 minutes). In one example, the values produced by a search query for a KPI by executing the search query on a schedule are stored in a data store, and are used to calculate an aggregate KPI score for a service, as described in greater detail below in conjunction with FIGS. **32-33**. An aggregate KPI score for the service is indicative of an overall performance of the KPIs of the service.

Referring to FIG. **24**, GUI **2400** can include a button **2416** to receive input specifying a frequency of monitoring (schedule) for determining the value produced by the search query of the KPI. The frequency of monitoring (e.g., schedule) of the KPI can be used to determine a resolution for an aggregate KPI score for the service. The aggregate KPI score for the service is indicative of an overall performance of the KPIs of the service. The accuracy of the aggregate KPI score for the service for a given point in time can be based on the frequency of monitoring of the KPI. For example, a higher frequency can provide higher resolution which can help produce a more accurate aggregate KPI score.

The machine data used by a search query defining a KPI to produce a value can be based on a time range. The time range can be a user-defined time range or a default time range. For example, in the service-monitoring dashboard example above, a user can select, via the service-monitoring dashboard, a time range to use (e.g., Last 15 minutes) to further specify, for example, based on time-stamps, which machine data should be used by a search query defining a KPI. In another example, the time range may be to use the machine data since the last time the value was produced by the search query. For example, if the KPI is assigned a frequency of monitoring of 5 minutes, then the search query can execute every 5 minutes, and for each execution use the machine data for the last 5 minutes relative to the execution time. In another implementation, the time range is a selected (e.g., user-selected) point in time and the definition of an individual KPI can specify the aggregation period for the respective KPI. By including the aggregation period for an individual KPI as part of the definition of the respective KPI, multiple KPIs can run on different aggregation periods, which can more accurately represent certain types of aggregations, such as, distinct counts and sums, improving the utility of defined thresholds.



In this manner, the value of each KPI can be displayed at a given point in time. In one example, a user may also select “real time” as the point in time to produce the most up to date value for each KPI using its respective individually defined aggregation period.

GUI **2400** can include a button **2414** to receive input assigning a weight to the KPI to indicate an importance of the KPI for the service relative to other KPIs defined for the service. The importance (e.g., weight) of the KPI can be used to determine an aggregate KPI score for the service, which is indicative of an overall performance of the KPIs of the service. Some implementations for using the importance and frequency of monitoring for each KPI to determine an aggregate KPI score for the service are discussed in greater detail below in conjunction with FIGS. **32-33**. FIG. **27** illustrates an example of a GUI **2700** of a service monitoring system for editing which entity definitions to use for a KPI, in accordance with one or more implementations of the present disclosure. GUI **2700** may be displayed in response to the user activation of button **2412** in GUI **2400** of FIG. **24**. GUI **2700** can include a button **2710** for creating a new entity definition. If button **2710** is selected, GUI **1600** in FIG. **16** can be displayed and an entity definition can be created as described above in conjunction with FIG. **6** and FIG. **16**.

Referring to FIG. **27**, GUI **2700** can display buttons **2701**, **2703** for receiving a selection of whether to include all of the entity definitions, which are associated with the service via the service definition, for the KPI. If the Yes button **2701** is selected, the search query for the KPI can produce a value derived from the machine data pertaining to all of the entities represented by the entity definitions that are included in the service definition for the service. If the No button **2703** is selected, a member list **2704** is displayed. The member list **2704** includes the entity definitions that are included in the service definition for the service. GUI **2700** can include a filter box **2702** to receive input for filtering the member list **2704** of entity definitions to display a subset of the entity definitions.

GUI **2700** can facilitate user input for selecting one or more entity definitions from the member list **2704** and dragging the selected entity definition(s) to an exclusion list **2712** to indicate that the entities identified in each selected entity definition should not be considered for the current KPI. This exclusion means that the search criteria of the search query defining the KPI is changed to no longer search for machine data pertaining to the entities identified in the entity definitions from the exclusion list **2712**. For example, entity definition **2705** (e.g., webserver07.splunk.com) can be selected and dragged to the exclusion list **2712**. When the search query for the KPI produces a value, the value will be derived from machine data, which does not include machine data pertaining to webserver07.splunk.com.

FIG. **28** is a flow diagram of an implementation of a method **2800** for defining one or more thresholds for a KPI, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block **2802**, the computing machine identifies a service definition for a service. In one implementation, the computing machine receives input (e.g., user input) selecting a ser-

vice definition. The computing machine accesses the service definition for a service from memory.

At block **2804**, the computing machine identifies a KPI for the service. In one implementation, the computing machine receives input (e.g., user input) selecting a KPI of the service. The computing machine accesses data representing the KPI from memory.

At block **2806**, the computing machine causes display of one or more graphical interfaces enabling a user to set a threshold for the KPI. The KPI can be in one of multiple states. Example states can include, and are not limited to, unknown, trivial state, informational state, normal state, warning state, error state, and critical state. Each state can be represented by a range of values. At a certain time, the KPI can be in one of the states depending on which range the value, which is produced by the search query for the KPI, falls into. Each threshold defines an end of a range of values, which represents one of the states. Some examples of graphical interfaces for enabling a user to set a threshold for the KPI are discussed in greater detail below in conjunction with FIG. **29A** to FIG. **31C**.

At block **2808**, the computing machine receives, through the graphical interfaces, an indication of how to set the threshold for the KPI. The computing machine can receive input (e.g., user input), via the graphical interfaces, specifying the field or alias that should be used for the threshold(s) for the KPI. The computing machine can also receive input (e.g., user input), via the graphical interfaces, of the parameters for each state. The parameters for each state can include, for example, and not limited to, a threshold that defines an end of a range of values for the state, a unique name, and one or more visual indicators to represent the state.

In one implementation, the computing machine receives input (e.g., user input), via the graphical interfaces, to set a threshold and to apply the threshold to the KPI as determined using the machine data from the aggregate of the entities associated with the KPI.

In another implementation, the computing machine receives input (e.g., user input), via the graphical interfaces, to set a threshold and to apply the threshold to a KPI as the KPI is determined using machine data on a per entity basis for the entities associated with the KPI. For example, the computing machine can receive a selection (e.g., user selection) to apply thresholds on a per entity basis, and the computing machine can apply the thresholds to the value of the KPI as the value is calculated per entity.

For example, the computing machine may receive input (e.g., user input), via the graphical interfaces, to set a threshold of being equal or greater than 80% for the KPI for Avg CPU Load, and the KPI is associated with three entities (e.g., Entity-1, Entity-2, and Entity-3). When the KPI is determined using data for Entity-1, the value for the KPI for Avg CPU Load may be at 50%. When the KPI is determined using data for Entity-2, the value for the KPI for Avg CPU Load may be at 50%. When the KPI is determined using data for Entity-3, the value for the KPI for Avg CPU Load may be at 80%. If the threshold is applied to the values of the aggregate of the entities (two at 50% and one at 80%), the aggregate value of the entities is 60%, and the KPI would not exceed the 80% threshold. If the threshold is applied using an entity basis for the thresholds (applied to the individual KPI values as calculated pertaining to each entity), the computing machine can determine that the KPI pertaining to one of the entities (e.g., Entity-3) satisfies the threshold by being equal to 80%.

At block **2810**, the computing machine determines whether to set another threshold for the KPI. The computing machine can receive input, via the graphical interface, indi-



31

cating there is another threshold to set for the KPI. If there is another threshold to set for the KPI, the computing machine returns to block **2808** to set the other threshold.

If there is not another threshold to set for the KPI (block **2810**), the computing machine determines whether to set a threshold for another KPI for the service at block **2812**. The computing machine can receive input, via the graphical interface, indicating there is a threshold to set for another KPI for the service. In one implementation, there are a maximum number of thresholds that can be set for a KPI. In one implementation, a same number of states are to be set for the KPIs of a service. In one implementation, a same number of states are to be set for the KPIs of all services. The service monitoring system can be coupled to a data store that stores configuration data that specifies whether there is a maximum number of thresholds for a KPI and the value for the maximum number, whether a same number of states is to be set for the KPIs of a service and the value for the number of states, and whether a same number of states is to be set for the KPIs of all of the service and the value for the number of states. If there is a threshold to set for another KPI, the computing machine returns to block **2804** to identify the other KPI.

At block **2814**, the computing machine stores the one or more threshold settings for the one or more KPIs for the service. The computing machine associates the parameters for a state defined by a corresponding threshold in a data store that is coupled to the computing machine.

As will be discussed in more detail below, implementations of the present disclosure provide a service-monitoring dashboard that includes KPI widgets ("widgets") to visually represent KPIs of the service. A widget can be a Noel gauge, a spark line, a single value, or a trend indicator. A Noel gauge is indicator of measurement as described in greater detail below in conjunction with FIG. **40**. A widget of a KPI can present one or more values indicating how a respective service or an aspect of a service is performing at one or more points in time. The widget can also illustrate (e.g., using visual indicators such as color, shading, shape, pattern, trend compared to a different time range, etc.) the KPI's current state defined by one or more thresholds of the KPI.

FIGS. **29A-B** illustrate examples of a graphical interface enabling a user to set one or more thresholds for the KPI, in accordance with one or more implementations of the present disclosure.

FIG. **29A** illustrates an example GUI **2900** for receiving input for search processing language **2902** for defining a search query, in accordance with one or more implementations of the present disclosure. The KPI can be in one of multiple states (e.g., normal, warning, critical). Each state can be represented by a range of values. At a certain time, the KPI can be in one of the states depending on which range the value, which is produced by the search query for the KPI, falls into. GUI **2900** can display an input box **2904** for a field to which the threshold(s) can be applied. In particular, a threshold can be applied to the value produced by the search query defining the KPI. The value can be, for example, the field's value extracted from an event when the search query is executed, a statistic calculated based on one or more values of the field in one or more events located when the search query is executed, a count of events satisfying the search criteria that include a constraint for the field, etc. GUI **2900** may include the name **2904** of the actual field used in the search query or the name of an alias that defines a desired statistic or count to be produced by the search query. For example, the threshold may be applied to an average response time produced by the search query, and the average response time can be defined by the alias "rsp\_time" in the input box **2904**.

32

FIG. **29B** illustrates an example GUI **2950** for receiving input for selecting a data model for defining a search query, in accordance with one or more implementations of the present disclosure. GUI **2950** can be displayed if a KPI is defined using a data model.

GUI **2950** in FIG. **29B** can include a statistical function **2954** to be used for producing a value when executing the search query of the KPI. As shown, the statistical function **2954** is a count, and the resulting statistic (the count value) should be compared with one or more thresholds of the KPI. The GUI **2950** also includes a button **2956** for creating the threshold(s) for the KPI. When either button **2906** is selected from GUI **2900** or button **2956** is selected from GUI **2950**, GUI **3000** of FIG. **30** is displayed.

FIG. **30** illustrates an example GUI **3000** for enabling a user to set one or more thresholds for the KPI, in accordance with one or more implementations of the present disclosure. Each threshold for a KPI defines an end of a range of values, which represents one of the states. GUI **3000** can display a button **3002** for adding a threshold to the KPI. If button **3002** is selected, a GUI for facilitating user input for the parameters for the state associated with the threshold can be displayed, as discussed in greater detail below in conjunction with FIGS. **31A-C**.

Referring to FIG. **30**, if button **3002** is selected three times, there will be three thresholds for the KPI. Each threshold defines an end of a range of values, which represents one of the states. GUI **3000** can display a UI element (e.g., column **3006**) that includes sections representing the defined states for the KPI, as described in greater detail below in conjunction with FIGS. **31A-C**. GUI **3000** can facilitate user input to specify a maximum value **3004** and a minimum value **3008** for defining a scale for a widget that can be used to represent the KPI on the service-monitoring dashboard. Some implementations of widgets for representing KPIs are discussed in greater detail below in conjunction with FIGS. **40-42** and FIGS. **44-46**.

Referring to FIG. **30**, GUI **3000** can optionally include a button **3010** for receiving input indicating whether to apply the threshold(s) to the aggregate of the KPIs of the service or to the particular KPI. Some implementations for applying the threshold(s) to the aggregate of the KPIs of the service or to a particular KPI are discussed in greater detail below in conjunction with FIGS. **32-34**.

FIG. **31A** illustrates an example GUI **3100** for defining threshold settings for a KPI, in accordance with one or more implementations of the present disclosure. GUI **3100** is a modified view of GUI **3000**, which is provided once the user has requested to add several thresholds for a KPI via button **3002** of GUI **3000**. In particular, in response to the user request to add a threshold, GUI **3100** dynamically adds a GUI element in a designated area of GUI **3100**. A GUI element can be in the form of an input box divided into several portions to receive various user input and visually illustrate the received input. The GUI element can represent a specific state of the KPI. When multiple states are defined for the KPI, several GUI elements can be presented in the GUI **3100**. For example, the GUI elements can be presented as input boxes of the same size and with the same input fields, and those input boxes can be positioned horizontally, parallel to each other, and resemble individual records from the same table. Alternatively, other types of GUI elements can be provided to represent the states of the KPI.

Each state of the KPI can have a name, and can be represented by a range of values, and a visual indicator. The range of values is defined by one or more thresholds that can provide the minimum end and/or the maximum end of the range of

33

values for the state. The characteristics of the state (e.g., the name, the range of values, and a visual indicator) can be edited via input fields of the respective GUI element.

In the example shown in FIG. 31A, GUI 3100 includes three GUI elements representing three different states of the KPI based on three added thresholds. These states include states 3102, 3104, and 3106.

For each state, GUI 3100 can include a GUI element that displays a name (e.g., a unique name for that KPI) 3109, a threshold 3110, and a visual indicator 3112 (e.g., an icon having a distinct color for each state). The unique name 3109, a threshold 3110, and a visual indicator 3112 can be displayed based on user input received via the input fields of the respective GUI element. For example, the name "Normal" can be specified for state 3106, the name "Warning" can be specified for state 3104, the name "Critical" can be specified for state 3102.

The visual indicator 3112 can be, for example, an icon having a distinct visual characteristic such as a color, a pattern, a shade, a shape, or any combination of color, pattern, shade and shape, as well as any other visual characteristics. For each state, the GUI element can display a drop-down menu 3114, which when selected, displays a list of available visual characteristics. A user selection of a specific visual characteristic (e.g., a distinct color) can be received for each state.

For each state, input of a threshold value representing the minimum end of the range of values for the corresponding state of the KPI can be received via the threshold portion 3110 of the GUI element. The maximum end of the range of values for the corresponding state can be either a preset value or can be defined by (or based on) the threshold associated with the succeeding state of the KPI, where the threshold associated with the succeeding state is higher than the threshold associated with the state before it.

For example, for Normal state 3106, the threshold value 0 may be received to represent the minimum end of the range of KPI values for that state. The maximum end of the range of KPI values for the Normal state 3106 can be defined based on the threshold associated with the succeeding state (e.g., Warning state 3104) of the KPI. For example, the threshold value 50 may be received for the Warning state 3104 of the KPI. Accordingly, the maximum end of the range of KPI values for the Normal state 3106 can be set to a number immediately preceding the threshold value of 50 (e.g., it can be set to 49 if the values used to indicate the KPI state are integers).

The maximum end of the range of KPI values for the Warning state 3104 is defined based on the threshold associated with the succeeding state (e.g., Critical state 3102) of the KPI. For example, the threshold value 75 may be received for the Critical state 3102 of the KPI, which may cause the maximum end of the range of values for the Warning state 3104 to be set to 74. The maximum end of the range of values for the highest state (e.g., Critical state 3102) can be a preset value or an indefinite value.

When input is received for a threshold value for a corresponding state of the KPI and/or a visual characteristic for an icon of the corresponding state of the KPI, GUI 3100 reflects this input by dynamically modifying a visual appearance of a vertical UI element (e.g., column 3118) that includes sections that represent the defined states for the KPI. Specifically, the sizes (e.g., heights) of the sections can be adjusted to visually illustrate ranges of KPI values for the states of the KPI, and the threshold values can be visually represented as marks on the column 3118. In addition, the appearance of each section is modified based on the visual characteristic (e.g., color, pattern) selected by the user for each state via a drop-down

34

menu 3114. In some implementations, once the visual characteristic is selected for a specific state, it is also illustrated by modified appearance (e.g., modified color or pattern) of icon 3112 positioned next to a threshold value associated with that state.

For example, if the color green is selected for the Normal state 3106, a respective section of column 3118 can be displayed with the color green to represent the Normal state 3106. In another example, if the value 50 is received as input for the minimum end of a range of values for the Warning state 3104, a mark 3117 is placed on column 3118 to represent the value 50 in proportion to other marks and the overall height of the column 3118. As discussed above, the size (e.g., height) of each section of the UI element (e.g., column) 3118 is defined by the minimum end and the maximum end of the range of KPI values of the corresponding state.

In one implementation, GUI 3100 displays one or more pre-defined states for the KPI. Each predefined state is associated with at least one of a pre-defined unique name, a pre-defined value representing a minimum end of a range of values, or a predefined visual indicator. Each pre-defined state can be represented in GUI 3100 with corresponding GUI elements as described above.

GUI 3100 can facilitate user input to specify a maximum value 3116 and a minimum value 3120 for the combination of the KPI states to define a scale for a widget that represents the KPI. Some implementations of widgets for representing KPIs are discussed in greater detail below in conjunction with FIGS. 40-42 and FIGS. 44-46. GUI 3100 can display a button 3122 for receiving input indicating whether to apply the threshold(s) to the aggregate KPI of the service or to the particular KPI or both. The application of threshold(s) to the aggregate KPI of the service or to a particular KPI is discussed in more detail below in conjunction with FIG. 33.

FIGS. 31B-31C illustrate GUIs for defining threshold settings for a KPI, in accordance with an alternative implementation of the present disclosure. In GUI 3150 of FIG. 31B, adjacent to column 3118, a line chart 3152 is displayed. The line chart 3152 represents the KPI values for the current KPI over a period of time selected from drop down menu 3154. The KPI values are plotted over the period of time on a first horizontal axis and against a range of values set by the maximum value 3116 and minimum value 3120 on a second vertical axis. In one implementation when a mark 3156 is added to column 3118 indicating the end of a range of values for the a particular state a horizontal line 3158 is displayed along the length of line chart 3152. The horizontal line 3158 makes it easy to visually correlate the KPI values represented by line chart 3152 with the end of the range of values. For example, in FIG. 31B, with the "Critical" state having a range below 15 GB, the horizontal line 3158 indicates that the KPI values drop below the end of the range four different times. This may provide information to a user that the end of the range of values indicated by mark 3156 can be adjusted.

In GUI 3160 of FIG. 31C, the user has adjusted the position of mark 3156, thereby decreasing the end of the range of values for the "Critical" state to 10 GB. Horizontal line 3158 is also lowered to reflect the change. In one implementation, the user may click and drag mark 3156 down to the desired value. In another implementation, the user may type in the desired value. The user can tell that the KPI values now drop below the end of the only once, thereby limiting the number of alerts associated with the defined threshold.

Aggregate Key Performance Indicators

FIG. 32 is a flow diagram of an implementation of a method 3200 for calculating an aggregate KPI score for a service based on the KPIs for the service, in accordance with one or

35

more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block **3201**, the computing machine identifies a service to evaluate. The service is provided by one or more entities. The computing system can receive user input, via one or more graphical interfaces, selecting a service to evaluate. The service can be represented by a service definition that associates the service with the entities as discussed in more detail above.

At block **3203**, the computing machine identifies key performance indicators (KPIs) for the service. The service definition representing the service can specify KPIs available for the service, and the computing machine can determine the KPIs for the service from the service definition of the service. Each KPI can pertain to a different aspect of the service. Each KPI can be defined by a search query that derives a value for that KPI from machine data pertaining to entities providing the service. As discussed above, the entities providing the service are identified in the service definition of the service. According to a search query, a KPI value can be derived from machine data of all or some entities providing the service.

In some implementations, not all of the KPIs for a service are used to calculate the aggregate KPI score for the service. For example, a KPI may solely be used for troubleshooting and/or experimental purposes and may not necessarily contribute to providing the service or impacting the performance of the service. The troubleshooting/experimental KPI can be excluded from the calculation of the aggregate KPI score for the service.

In one implementation, the computing machine uses a frequency of monitoring that is assigned to a KPI to determine whether to include a KPI in the calculation of the aggregate KPI score. The frequency of monitoring is a schedule for executing the search query that defines a respective KPI. As discussed above, the individual KPIs can represent saved searches. These saved searches can be scheduled for execution based on the frequency of monitoring of the respective KPIs. In one example, the frequency of monitoring specifies a time period (e.g., 1 second, 2 minutes, 10 minutes, 30 minutes, etc.) for executing the search query that defines a respective KPI, which then produces a value for the respective KPI with each execution of the search query. In another example, the frequency of monitoring specifies particular times (e.g., 6:00 am, 12:00 pm, 6:00 pm, etc.) for executing the search query. The values produced for the KPIs of the service, based on the frequency of monitoring for the KPIs, can be considered when calculating a score for an aggregate KPI of the service, as discussed in greater detail below in conjunction with FIG. **34**.

Alternatively, the frequency of monitoring can specify that the KPI is not to be measured (that the search query for a KPI is not to be executed). For example, a troubleshooting KPI may be assigned a frequency of monitoring of zero.

In one implementation, if a frequency of monitoring is unassigned for a KPI, the KPI is automatically excluded in the calculation for the aggregate KPI score. In one implementation, if a frequency of monitoring is unassigned for a KPI, the KPI is automatically included in the calculation for the aggregate KPI score.

The frequency of monitoring can be assigned to a KPI automatically (without any user input) based on default set-

36

tings or based on specific characteristics of the KPI such as a service aspect associated with the KPI, a statistical function used to derive a KPI value (e.g., maximum versus average), etc. For example, different aspects of the service can be associated with different frequencies of monitoring, and KPIs can inherit frequencies of monitoring of corresponding aspects of the service.

Values for KPIs can be derived from machine data that is produced by different sources. The sources may produce the machine data at various frequencies (e.g., every minute, every 10 minutes, every 30 minutes, etc.) and/or the machine data may be collected at various frequencies (e.g., every minute, every 10 minutes, every 30 minutes, etc.). In another example, the frequency of monitoring can be assigned to a KPI automatically (without any user input) based on the accessibility of machine data associated with the KPI (associated through entities providing the service). For example, an entity may be associated with machine data that is generated at a medium frequency (e.g., every 10 minutes), and the KPI for which a value is being produced using this particular machine data can be automatically assigned a medium frequency for its frequency of monitoring.

Alternatively, frequency of monitoring can be assigned to KPIs based on user input. FIG. **33A** illustrates an example GUI **3300** for creating and/or editing a KPI, including assigning a frequency of monitoring to a KPI, based on user input, in accordance with one or more implementations of the present disclosure. GUI **3300** for can include a button **3311** to receive a user request to assign a frequency of monitoring to the KPI being created or modified. Upon activating button **3311**, a user can enter (e.g., via another GUI or a command line interface) a frequency (e.g., a user defined value) for the KPI, or select a frequency from a list presented to the user. In one example, the list may include various frequency types, where each frequency type is mapped to a pre-defined and/or user-defined time period. For example, the frequency types may include Real Time (e.g., 1 second), High Frequency (e.g., 2 minutes), Medium Frequency (e.g., 10 minutes), Low Frequency (e.g., 30 minutes), Do Not Measure (e.g., no frequency).

The assigned frequency of monitoring of KPIs can be included in the service definition specifying the KPIs, or in a separate data structure together with other settings of a KPI.

Referring to FIG. **32**, at block **3205**, the computing machine derives one or more values for each of the identified KPIs. The computing machine can cause the search query for each KPI to execute to produce a corresponding value. In one implementation, as discussed above, the search query for a particular KPI is executed based on a frequency of monitoring assigned to the particular KPI. When the frequency of monitoring for a KPI is set to a time period, for example, High Frequency (e.g., 2 minutes), a value for the KPI is derived each time the search query defining the KPI is executed every 2 minutes. The derived value(s) for each KPI can be stored in an index. In one implementation, when a KPI is assigned a frequency of monitoring of Do Not Measure or is assigned a zero frequency (no frequency), no value is produced (the search query for the KPI is not executed) for the respective KPI and no values for the respective KPI are stored in the data store.

At block **3207**, the computing machine calculates a value for an aggregate KPI score for the service using the value(s) from each of the KPIs of the service. The value for the aggregate KPI score indicates an overall performance of the service. For example, a Web Hosting service may have 10 KPIs and one of the 10 KPIs may have a frequency of monitoring set to Do Not Monitor. The other nine KPIs may be assigned

various frequencies of monitoring. The computing machine can access the values produced for the nine KPIs in the data store to calculate the value for the aggregate KPI score for the service, as discussed in greater detail below in conjunction with FIG. 34. Based on the values obtained from the data store, if the values produced by the search queries for 8 of the 9 KPIs indicate that the corresponding KPI is in a normal state, then the value for an aggregate KPI score may indicate that the overall performance of the service is normal.

An aggregate KPI score can be calculated by adding the values of all KPIs of the same service together. Alternatively, an importance of each individual KPI relative to other KPIs of the service is considered when calculating the aggregate KPI score for the service. For example, a KPI can be considered more important than other KPIs of the service if it has a higher importance weight than the other KPIs of the service.

In some implementations, importance weights can be assigned to KPIs automatically (without any user input) based on characteristics of individual KPIs. For example, different aspects of the service can be associated with different weights, and KPIs can inherit weights of corresponding aspects of the service. In another example, a KPI deriving its value from machine data pertaining to a single entity can be automatically assigned a lower weight than a KPI deriving its value from machine data pertaining to multiple entities, etc.

Alternatively, importance weights can be assigned to KPIs based on user input. Referring again to FIG. 33A, GUI 3300 can include a button 3309 to receive a user request to assign a weight to the KPI being created or modified. Upon selecting button 3309, a user can enter (e.g., via another GUI or a command line interface) a weight (e.g., a user defined value) for the KPI, or select a weight from a list presented to the user. In one implementation, a greater value indicates that a greater importance is placed on a KPI. For example, the set of values may be 1-10, where the value 10 indicates high importance of the KPI relative to the other KPIs for the service. For example, a Web Hosting service may have three KPIs: (1) CPU Usage, (2) Memory Usage, and (3) Request Response Time. A user may provide input indicating that the Request Response Time KPI is the most important KPI and may assign a weight of 10 to the Request Response Time KPI. The user may provide input indicating that the CPU Usage KPI is the next most important KPI and may assign a weight of 5 to the CPU Usage KPI. The user may provide input indicating that the Memory Usage KPI is the least important KPI and may assign a weight of 1 to the Memory Usage KPI.

In one implementation, a KPI is assigned an overriding weight. The overriding weight is a weight that overrides the importance weights of the other KPIs of the service. Input (e.g., user input) can be received for assigning an overriding weight to a KPI. The overriding weight indicates that the status (state) of KPI should be used a minimum overall state of the service. For example, if the state of the KPI, which has the overriding weight, is warning, and one or more other KPIs of the service have a normal state, then the service may only be considered in either a warning or critical state, and the normal state(s) for the other KPIs can be disregarded.

In another example, a user can provide input that ranks the KPIs of a service from least important to most important, and the ranking of a KPI specifies the user selected weight for the respective KPI. For example, a user may assign a weight of 1 to the Memory Usage KPI, assign a weight of 2 to the CPU Usage KPI, and assign a weight of 3 to the Request Response Time KPI. The assigned weight of each KPI may be included in the service definition specifying the KPIs, or in a separate data structure together with other settings of a KPI.

Alternatively or in addition, a KPI can be considered more important than other KPIs of the service if it is measured more frequently than the other KPIs of the service. In other words, search queries of different KPIs of the service can be executed with different frequency (as specified by a respective frequency of monitoring) and queries of more important KPIs can be executed more frequently than queries of less important KPIs.

As will be discussed in more detail below in conjunction with FIG. 34, the calculation of a score for an aggregate KPI may be based on ratings assigned to different states of an individual KPI. Referring again to FIG. 33A, a user can select button 3313 for defining threshold settings, including state ratings, for a KPI to display GUI 3350 in FIG. 33B. FIG. 33B illustrates an example GUI 3350 for defining threshold settings, including state ratings, for a KPI, in accordance with one or more implementations of the present disclosure. Similarly to GUI 3100 of FIG. 31A, GUI 3350 includes horizontal GUI elements (e.g., in the form of input boxes) 3352, 3354 and 3356 that represent specific states of the KPI. For each state, a corresponding GUI element can display a name 3359, a threshold 3360, and a visual indicator 3362 (e.g., an icon having a distinct color for each state). The name 3359, a threshold 3360, and a visual indicator 3362 can be displayed based on user input received via the input fields of the respective GUI element. GUI 3350 can include a vertical GUI element (e.g., a column) 3368 that changes appearance (e.g., the size and color of its sectors) based on input received for a threshold value for a corresponding state of the KPI and/or a visual characteristic for an icon of the corresponding state of the KPI. In some implementations, once the visual characteristic is selected for a specific state via the menu 3364, it is also illustrated by the modified appearance (e.g., modified color or pattern) of icon 3362 positioned next to a threshold value associated with that state.

In addition, GUI 3350 provides for configuring a rating for each state of the KPI. The ratings indicate which KPIs should be given more or less consideration in view of their current states. When calculating an aggregate KPI, a score of each individual KPI reflects the rating of that KPI's current state, as will be discussed in more detail below in conjunction with FIG. 34. Ratings for different KPI states can be assigned automatically (e.g., based on a range of KPI values for a state) or specified by a user. GUI 3350 can include a field 3380 that displays an automatically generated rating or a rating entered or selected by a user. Field 3380 may be located next to (or in the same row as) a horizontal GUI element representing a corresponding state. Alternatively, field 3380 can be part of the horizontal GUI element. In one example, a user may provide input assigning a rating of 1 to the Normal State, a rating of 2 to the Warning State, and a rating of 3 to the Critical State.

In one implementation, GUI 3350 displays a button 3372 for receiving input indicating whether to apply the threshold(s) to the aggregate KPI of the service or to the particular KPI or both. If a threshold is configured to be applied to a certain individual KPI, then a specified action (e.g., generate alert, add to report) will be triggered when a value of that KPI reaches (or exceeds) the individual KPI threshold. If a threshold is configured to be applied to the aggregate KPI of the service, then a specified action (e.g., create notable event, generate alert, add to incident report) will be triggered when a value (e.g., a score) of the aggregate KPI reaches (or exceeds) the aggregate KPI threshold. In some implementations, a threshold can be applied to both or either the individual or aggregate KPI, and different actions or the same action can be triggered depending on the KPI to which the

threshold is applied. The actions to be triggered can be pre-defined or specified by the user via a user interface (e.g., a GUI or a command line interface) while the user is defining thresholds or after the thresholds have been defined. The action to be triggered in view of thresholds can be included in the service definition identifying the respective KPI(s) or can be stored in a data structure dedicated to store various KPI settings of a relevant KPI.

FIG. 34 is a flow diagram of an implementation of a method 3400 for calculating a score for an aggregate KPI for the service, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block 3402, the computing machine identifies a service to be evaluated. The service is provided by one or more entities. The computing system can receive user input, via one or more graphical interfaces, selecting a service to evaluate.

At block 3404, the computing machine identifies key performance indicators (KPIs) for the service. The computing machine can determine the KPIs for the service from the service definition of the service. Each KPI indicates how a specific aspect of the service is performing at a point in time.

At block 3406, the computing machine optionally identifies a weighting (e.g., user selected weighting or automatically assigned weighting) for each of the KPIs of the service. As discussed above, the weighting of each KPI can be determined from the service definition of the service or a KPI definition storing various setting of the KPI.

At block 3408, the computing machine derives one or more values for each KPI for the service by executing a search query associated with the KPI. As discussed above, each KPI is defined by a search query that derives the value for a corresponding KPI from the machine data that is associated with the one or more entities that provide the service.

As discussed above, the machine data associated with the one or more entities that provide the same service is identified using a user-created service definition that identifies the one or more entities that provide the service. The user-created service definition also identifies, for each entity, identifying information for locating the machine data pertaining to that entity. In another example, the user-created service definition also identifies, for each entity, identifying information for a user-created entity definition that indicates how to locate the machine data pertaining to that entity. The machine data can include for example, and is not limited to, unstructured data, log data, and wire data. The machine data associated with an entity can be produced by that entity. In addition or alternatively, the machine data associated with an entity can include data about the entity, which can be collected through an API for software that monitors that entity.

The computing machine can cause the search query for each KPI to execute to produce a corresponding value for a respective KPI. The search query defining a KPI can derive the value for that KPI in part by applying a late-binding schema to machine data or, more specifically, to events containing raw portions of the machine data. The search query can derive the value for the KPI by using a late-binding schema to extract an initial value and then performing a calculation on (e.g., applying a statistical function to) the initial value.

The values of each of the KPIs can differ at different points in time. As discussed above, the search query for a KPI can be executed based on a frequency of monitoring assigned to the particular KPI. When the frequency of monitoring for a KPI is set to a time period, for example, Medium Frequency (e.g., 10 minutes), a value for the KPI is derived each time the search query defining the KPI is executed every 10 minutes. The derived value(s) for each KPI can be stored in a data store. When a KPI is assigned a zero frequency (no frequency), no value is produced (the search query for the KPI is not executed) for the respective KPI.

The derived value(s) of a KPI is indicative of how an aspect of the service is performing. In one example, the search query can derive the value for the KPI by applying a late-binding schema to machine data pertaining to events to extract values for a specific fields defined by the schema. In another example, the search query can derive the value for that KPI by applying a late-binding schema to machine data pertaining to events to extract an initial value for a specific field defined by the schema and then performing a calculation on (e.g., applying a statistical function to) the initial value to produce the calculation result as the KPI value. In yet another example, the search query can derive the value for the KPI by applying a late-binding schema to machine data pertaining to events to extract an initial value for specific fields defined by the late-binding schema to find events that have certain values corresponding to the specific fields, and counting the number of found events to produce the resulting number as the KPI value.

At block 3410, the computing machine optionally maps the value produced by a search query for each KPI to a state. As discussed above, each KPI can have one or more states defined by one or more thresholds. In particular, each threshold can define an end of a range of values. Each range of values represents a state for the KPI. At a certain point in time or a period of time, the KPI can be in one of the states (e.g., normal state, warning state, critical state) depending on which range the value, which is produced by the search query of the KPI, falls into. For example, the value produced by the Memory Usage KPI may be in the range representing a Warning State. The value produced by the CPU Usage KPI may be in the range representing a Warning State. The value produced by the Request Response Time KPI may be in the range representing a Critical State.

At block 3412, the computing machine optionally maps the state for each KPI to a rating assigned to that particular state for a respective KPI (e.g., automatically or based on user input). For example, for a particular KPI, a user may provide input assigning a rating of 1 to the Normal State, a rating of 2 to the Warning State, and a rating of 3 to the Critical State. In some implementations, the same ratings are assigned to the same states across the KPIs for a service. For example, the Memory Usage KPI, CPU Usage KPI, and Request Response Time KPI for a Web Hosting service may each have Normal State with a rating of 1, a Warning State with a rating of 2, and a Critical State with a rating of 3. The computing machine can map the current state for each KPI, as defined by the KPI value produced by the search query, to the appropriate rating. For example, the Memory Usage KPI in the Warning State can be mapped to 2. The CPU Usage KPI in the Warning State can be mapped to 2. The Request Response Time KPI in the Critical State can be mapped to 3. In some implementations, different ratings are assigned to the same states across the KPIs for a service. For example, the Memory Usage KPI may each have Critical State with a rating of 3, and the Request Response Time KPI may have Critical State with a rating of 5.

At block 3414, the computing machine calculates an impact score for each KPI. In some implementations, the impact score of each KPI can be based on the importance weight of a corresponding KPI (e.g.,  $\text{weight} \times \text{KPI value}$ ). In other implementations, the impact score of each KPI can be based on the rating associated with a current state of a corresponding KPI (e.g.,  $\text{rating} \times \text{KPI value}$ ). In yet other implementations, the impact score of each KPI can be based on both the importance weight of a corresponding KPI and the rating associated with a current state of the corresponding KPI. For example, the computing machine can apply the weight of the KPI to the rating for the state of the KPI. The impact of a particular KPI at a particular point in time on the aggregate KPI can be the product of the rating of the state of the KPI and the importance (weight) assigned to the KPI. In one implementation, the impact score of a KPI can be calculated as follows:

$$\text{Impact Score of KPI} = (\text{weight}) \times (\text{rating of state})$$

For example, when the weight assigned to the Memory Usage KPI is 1 and the Memory Usage KPI is in a Warning State, the impact score of the Memory Usage KPI =  $1 \times 2$ . When the weight assigned to the CPU Usage KPI is 2 and the CPU Usage KPI is in a Warning State, the impact score of the CPU Usage KPI =  $2 \times 2$ . When the weight assigned to the Request Response Time KPI is 3 and the Request Response Time KPI is in a Critical State, the impact score of the Request Response Time KPI =  $3 \times 3$ .

In another implementation, the impact score of a KPI can be calculated as follows:

$$\text{Impact Score of KPI} = (\text{weight}) \times (\text{rating of state}) \times (\text{value})$$

In yet some implementations, the impact score of a KPI can be calculated as follows:

$$\text{Impact Score of KPI} = (\text{weight}) \times (\text{value})$$

At block 3416, the computing machine calculates an aggregate KPI score ("score") for the service based on the impact scores of individual KPIs of the service. The score for the aggregate KPI indicates an overall performance of the service. The score of the aggregate KPI can be calculated periodically (as configured by a user or based on a default time interval) and can change over time based on the performance of different aspects of the service at different points in time. For example, the aggregate KPI score may be calculated in real time (continuously calculated until interrupted). The aggregate KPI score may be calculated may be calculated periodically (e.g., every second).

In some implementations, the score for the aggregate KPI can be determined as the sum of the individual impact scores for the KPIs of the service. In one example, the aggregate KPI score for the Web Hosting service can be as follows:

$$\begin{aligned} \text{Aggregate KPI}_{\text{Web Hosting}} &= (\text{weight} \times \text{rating of state})_{\text{Memory Usage KPI}} + (\text{weight} \times \text{rating of state})_{\text{CPU Usage KPI}} + (\text{weight} \times \text{rating of state})_{\text{Request Response Time KPI}} \\ &= (1 \times 2) + (2 \times 2) + (3 \times 3) = 15. \end{aligned}$$

In another example, the aggregate KPI score for the Web Hosting service can be as follows:

$$\begin{aligned} \text{Aggregate KPI}_{\text{Web Hosting}} &= (\text{weight} \times \text{rating of state} \times \text{value})_{\text{Memory Usage KPI}} + (\text{weight} \times \text{rating of state} \times \text{value})_{\text{CPU Usage KPI}} + (\text{weight} \times \text{rating of state} \times \text{value})_{\text{Request Response Time KPI}} \\ &= (1 \times 2 \times 60) + (2 \times 2 \times 55) + (3 \times 3 \times 80) = 1060. \end{aligned}$$

In yet some other implementations, the impact score of an aggregate KPI can be calculated as a weighted average as follows:

$$\begin{aligned} \text{Aggregate KPI}_{\text{Web Hosting}} &= [(\text{weight} \times \text{rating of state})_{\text{Memory Usage KPI}} + (\text{weight} \times \text{rating of state})_{\text{CPU Usage KPI}} + (\text{weight} \times \text{rating of state})_{\text{Request Response Time KPI}}] / \\ &= (\text{weight}_{\text{Memory Usage KPI}} + \text{weight}_{\text{CPU Usage KPI}} + \text{weight}_{\text{Request Response Time KPI}}) \end{aligned}$$

A KPI can have multiple values produced for the particular KPI for different points in time, for example, as specified by a frequency of monitoring for the particular KPI. The multiple values for a KPI can be that in a data store. In one implementation, the latest value that is produced for the KPI is used for calculating the aggregate KPI score for the service, and the individual impact scores used in the calculation of the aggregate KPI score can be the most recent impact scores of the individual KPIs based on the most recent values for the particular KPI stored in a data store. Alternatively, a statistical function (e.g., average, maximum, minimum, etc.) is performed on the set of the values that is produced for the KPI is used for calculating the aggregate KPI score for the service. The set of values can include the values over a time period between the last calculation of the aggregate KPI score and the present calculation of the aggregate KPI score. The individual impact scores used in the calculation of the aggregate KPI score can be average impact scores, maximum impact score, minimum impact scores, etc. over a time period between the last calculation of the aggregate KPI score and the present calculation of the aggregate KPI score.

The individual impact scores for the KPIs can be calculated over a time range (since the last time the KPI was calculated for the aggregate KPI score). For example, for a Web Hosting service, the Request Response Time KPI may have a high frequency (e.g., every 2 minutes), the CPU Usage KPI may have a medium frequency (e.g., every 10 minutes), and the Memory Usage KPI may have a low frequency (e.g., every 30 minutes). That is, the value for the Memory Usage KPI can be produced every 30 minutes using machine data received by the system over the last 30 minutes, the value for the CPU Usage KPI can be produced every 10 minutes using machine data received by the system over the last 10 minutes, and the value for the Request Response Time KPI can be produced every 2 minutes using machine data received by the system over the last 2 minutes. Depending on the point in time for when the aggregate KPI score is being calculated, the value (e.g., and thus state) of the Memory Usage KPI may not have been refreshed (the value is stale) because the Memory Usage KPI has a low frequency (e.g., every 30 minutes). Whereas, the value (e.g., and thus state) of the Request Response Time KPI used to calculate the aggregate KPI score is more likely to be refreshed (reflect a more current state) because the Request Response Time KPI has a high frequency (e.g., every 2 minutes). Accordingly, some KPIs may have more impact on how the score of the aggregate KPI changes overtime than other KPIs, depending on the frequency of monitoring of each KPI.

In one implementation, the computing machine causes the display of the calculated aggregate KPI score in one or more graphical interfaces and the aggregate KPI score is updated in the one or more graphical interfaces each time the aggregate KPI score is calculated. In one implementation, the configuration for displaying the calculated aggregate KPI in one or more graphical interfaces is received as input (e.g., user input), stored in a data store coupled to the computing machine, and accessed by the computing machine.

At block 3418, the computing machine compares the score for the aggregate KPI to one or more thresholds. As discussed above with respect to FIG. 33B, one or more thresholds can be defined and can be configured to apply to a specific individual KPI and/or an aggregate KPI including the specific individual

KPI. The thresholds can be stored in a data store that is coupled to the computing machine. If the thresholds are configured to be applied to the aggregate KPI, the computing machine compares the score of the aggregate KPI to the thresholds. If the computing machine determines that the aggregate KPI score exceeds or reaches any of the thresholds, the computing machine determines what action should be triggered in response to this comparison.

Referring to FIG. 34, at block 3420, the computing machine causes an action to be performed based on the comparison of the aggregate KPI score with the one or more thresholds. For example, the computing machine can generate an alert if the aggregate KPI score exceeds or reaches a particular threshold (e.g., the highest threshold). In another example, the computing machine can generate a notable event if the aggregate KPI score exceeds or reaches a particular threshold (e.g., the second highest threshold). In one implementation, the KPIs of multiple services is aggregated and used to create a notable event. In one implementation, the configuration for which of one or more actions to be performed is received as input (e.g., user input), stored in a data store coupled to the computing machine, and accessed by the computing machine.

#### Correlation Search and KPI Distribution Thresholding

As discussed above, the aggregate KPI score a service can be used to generate notable events and/or alarms, according to one or more implementations of the present disclosure. In another implementation, a correlation search is created and used to generate notable event(s) and/or alarm(s). A correlation search can be created to determine the status of a set of KPIs for a service over a defined window of time. Thresholds can be set on the distribution of the state of each individual KPI and if the distribution thresholds are exceeded then an alert/alarm can be generated.

The correlation search can be based on a discrete mathematical calculation. For example, the correlation search can include, for each KPI included in the correlation search, the following:

```
(sum_crit>threshold_crit)&&((sum_crit+sum_warn)>
(threshold_crit+threshold_warn))&&((sum_crit+
sum_warn+sum_normal)>(threshold_crit+threshold_warn+threshold_normal))
```

Input (e.g., user input) can be received that defines one or more thresholds for the counts of each state in a defined (e.g., user-defined) time window for each KPI. The thresholds define a distribution for the respective KPI. The distribution shift between states for the respective KPI can be determined. When the distribution for a respective KPI shifts toward a particular state (e.g., critical state), the KPI can be categorized accordingly. The distribution shift for each KPI can be determined, and each KPI can be categorized accordingly. When the KPIs for a service are categorized, the categorized KPIs can be compared to criteria for triggering a notable event. If the criteria are satisfied, a notable event can be triggered.

For example, a Web Hosting service may have three KPIs: (1) CPU Usage, (2) Memory Usage, and (3) Request Response Time. The counts for each state in a defined (e.g., user-defined) time window for the CPU Usage KPI can be determined, and the distribution thresholds can be applied to the counts. The distribution for the CPU Usage KPI may shift towards a critical state, and the CPU Usage KPI is flagged as critical accordingly. The counts for each state in a defined time window for the Memory Usage KPI can be determined, and the distribution thresholds for the Memory Usage KPI may also shift towards a critical state, and the Memory Usage KPI is flagged as critical accordingly.

The counts of each state in a defined time window for the Request Response Time KPI can be determined, and the distribution thresholds for the Request Response Time KPI can be applied to the counts. The distribution for the Request Response Time KPI may also shift towards a critical state, and the Request Response Time KPI is flagged as critical accordingly. The categories for the KPIs can be compared to the one or more criteria for triggering a notable event, and a notable event is triggered as a result of each of the CPU Usage KPI, Memory Usage KPI, and Request Response Time KPI being flagged as critical.

Input (e.g., user input) can be received specifying one or more criteria for triggering a notable event. For example, the criteria may be that when all of the KPIs in the correlation search for a service are flagged (categorized) a critical state, a notable event is triggered. In another example, the criteria may be that when a particular KPIs is flagged a particular state for a particular number of times, a notable event is triggered. Each KPI can be assigned a set of criteria.

For example, a Web Hosting service may have three KPIs: (1) CPU Usage, (2) Memory Usage, and (3) Request Response Time. The counts of each state in a defined (e.g., user-defined) time window for the CPU Usage KPI can be determined, and the distribution thresholds can be applied to the counts. The distribution for the CPU Usage KPI may shift towards a critical state, and the CPU Usage KPI is flagged as critical accordingly. The counts of each state in a defined time window for the Memory Usage KPI can be determined, and the distribution thresholds for the Memory Usage KPI can be applied to the counts. The distribution for the Memory Usage KPI may also shift towards a critical state, and the Memory Usage KPI is flagged as critical accordingly. The counts of each state in a defined time window for the Request Response Time KPI can be determined, and the distribution thresholds for the Request Response Time KPI can be applied to the counts. The distribution for the Request Response Time KPI may also shift towards a critical state, and the Request Response Time KPI is flagged as critical accordingly. The categories for the KPIs can be compared to the one or more criteria for triggering a notable event, and a notable event is triggered as a result of each of the CPU Usage KPI, Memory Usage KPI, and Request Response Time KPI being flagged as critical.

#### Example Service-Monitoring Dashboard

FIG. 35 is a flow diagram of an implementation of a method 3500 for creating a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block 3501, the computing machine causes display of a dashboard-creation graphical interface that includes a modifiable dashboard template, and a KPI-selection interface. A modifiable dashboard template is part of a graphical interface to receive input for editing/creating a custom service-monitoring dashboard. A modifiable dashboard template is described in greater detail below in conjunction with FIG. 36B. The display of the dashboard-creation graphical interface can be caused, for example, by a user selecting to create a service-monitoring dashboard from a GUI. FIG. 36A illustrates an example GUI 3650 for creating and/or editing a service-monitoring dashboard, in accordance with one or



45

more implementations of the present disclosure. In one implementation, GUI 3650 includes a menu item, such as Service-Monitoring Dashboards 3652, which when selected can present a list 3656 of existing service-monitoring dashboards that have already been created. The list 3656 can represent service-monitoring dashboards that have data that is stored in a data store for displaying the service-monitoring dashboards. Each service-monitoring dashboard in the list 3656 can include a button 3658 for requesting a drop-down menu listing editing options to edit the corresponding service-monitoring dashboard. Editing can include editing the service-monitoring dashboard and/or deleting the service-monitoring dashboard. When an editing option is selected from the drop-down menu, one or more additional GUIs can be displayed for editing the service-monitoring dashboard.

The dashboard creation graphical interface can be a wizard or any other type of tool for creating a service-monitoring dashboard that presents a visual overview of how one or more services and/or one or more aspects of the services are performing. The services can be part of an IT environment and can include, for example, a web hosting service, an email service, a database service, a revision control service, a sandbox service, a networking service, etc. A service can be provided by one or more entities such as host machines, virtual machines, switches, firewalls, routers, sensors, etc. Each entity can be associated with machine data that can have different formats and/or use different aliases for the entity. As discussed above, each service can be associated with one or more KPIs indicating how aspects of the service are performing. The KPI-selection interface of the dashboard creation GUI allows a user to select KPIs for monitoring the performance of one or more services, and the modifiable dashboard template of the dashboard creation GUI allows the user to specify how these KPIs should be presented on a service-monitoring dashboard that will be created based on the dashboard template. The dashboard template can also define the overall look of the service-monitoring dashboard. The dashboard template for the particular service-monitoring dashboard can be saved, and subsequently, the service-monitoring dashboard can be generated for display based on the customized dashboard template and KPI values derived from machine data, as will be discussed in more details below.

GUI 3650 can include a button 3654 that a user can activate to proceed to the creation of a service-monitoring dashboard, which can lead to GUI 3600 of FIG. 36B. FIG. 36B illustrates an example dashboard-creation GUI 3600 for creating a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure. GUI 3600 includes a modifiable dashboard template 3608 and a KPI-selection interface 3606 for selecting a key performance indicator (KPI) of a service. GUI 3600 can facilitate input (e.g., user input) of a name 3602 of the particular service-monitoring dashboard that is being created and/or edited. GUI 3600 can include a button 3612 for storing the dashboard template 3608 for creating the service-monitoring dashboard. GUI 3600 can display a set of identifiers 3604, each corresponding to a service. The set of identifiers 3604 is described in greater detail below. GUI 3600 can also include a configuration interface 3610 for configuring style settings pertaining to the service-monitoring dashboard. The configuration interface 3610 is described in greater detail below. GUI 3600 can also include a customization toolbar 3601 for customizing the service-monitoring dashboard as described in greater detail below in conjunction with FIG. 35. The configuration interface 3610 can also include entity identifiers and facilitate input (e.g., user input) for selecting entity identifier of entities to be included in the service-monitoring dashboard.

46

Returning to FIG. 35, at block 3503, the computing machine optionally receives, via the dashboard-creation graphical interface, input for customizing an image for the service-monitoring dashboard and causes the customized image to be displayed in the dashboard-creation graphical interface at block 3505. In one example, the computing machine optionally receives, via the dashboard-creation graphical interface, a selection of a background image for the service-monitoring dashboard and causes the selected background image to be displayed in the dashboard-creation graphical interface. The computing machine can display the selected background image in the modifiable dashboard template. FIG. 37 illustrates an example GUI 3700 for a dashboard-creation graphical interface including a user selected background image, in accordance with one or more implementations of the present disclosure. GUI 3700 displays the user selected image 3708 in the modifiable dashboard template 3710.

Referring again to FIG. 35, in another example, at block 3503, the computing machine optionally receives input (e.g., user input) via a customization toolbar (e.g., customization toolbar 3601 in FIG. 36B) for customizing an image for the service-monitoring dashboard. The customization toolbar can be a graphical interface containing drawing tools to customize a service-monitoring dashboard to define, for example, flow charts, text and connections between different elements on the service-monitoring dashboard. For example, the computing machine can receive input of a user drawing a flow chart or a representation of an environment (e.g., IT environment). In another example, the computing machine can receive input of a user drawing a representation of an entity and/or service. In another example, the computing machine can receive input of a user selection of an image to represent of an entity and/or service.

At block 3507, the computing machine receives, through the KPI-selection interface, a selection of a particular KPI for a service. As discussed above, each KPI indicates how an aspect of the service is performing at one or more points in time. A KPI is defined by a search query that derives one or more values for the KPI from the machine data associated with the one or more entities that provide the service whose performance is reflected by the KPI.

In one example, prior to receiving the selection of the particular KPI, the computing machine causes display of a context panel graphical interface in the dashboard-creation graphical interface that contains service identifiers for the services (e.g., all of the services) within an environment (e.g., IT environment). The computing machine can receive input, for example, of a user selecting one or more of the service identifiers, and dragging and placing one or more of the service identifiers on the dashboard template. In another example, the computing machine causes display of a search box to receive input for filtering the service identifiers for the services.

In another example, prior to receiving the selection of the particular KPI, the computing machine causes display of a drop-down menu of selectable services in the KPI selection interface, and receives a selection of one of the services from the drop-down menu. In some implementations, selectable services can be displayed as identifiers corresponding to individual services, where each identifier can be, for example, the name of a particular service or the name of a service definition representing the particular service. As discussed in more detail above, a service definition can associate the service with one or more entities (and thereby with heterogeneous machine data pertaining to the entities) providing the service,



and can specify one or more KPIs created for the service to monitor the performance of different aspects of the service.

In response to the user selection of a particular service, the computing machine can cause display of a list of KPIs associated with the selected service in the KPI selection interface, and can receive the user selection of the particular KPI from this list.

Referring again to FIG. 37, a user may select Web Hosting service 3701 in FIG. 37 from the set of KPI identifiers 3702, and in response to the selection of the Web Hosting service 3701, the computing machine can cause display of a set of KPIs available for the Web Hosting service 3701. FIG. 38 illustrates an example GUI 3800 for displaying a set of KPIs associated with a selected service, in accordance with one or more implementations of the present disclosure. GUI 3800 can be a pop-up window that includes a drop-down menu 3801, which when selected, displays a set of KPIs (e.g., Request Response Time and CPU Usage) associated with the service (e.g., Web Hosting service) corresponding to the selected service identifier. The user can then select a particular KPI from the menu. In another implementation, GUI 3800 also displays an aggregate KPI associated with the selected service, which can be selected to be represented by a KPI widget in the dashboard template for display in the service-monitoring dashboard.

Returning to FIG. 35, at block 3509, the computing machine receives a selection of a location for placing the selected KPI in the dashboard template for displaying a KPI widget in a dashboard. Each KPI widget can provide a numerical or graphical representation of one or more values for a corresponding KPI or service health score (aggregate KPI for a service) indicating how a service or an aspect of a service is performing at one or more points in time. For example, a user can select the desired location for a KPI widget by clicking (or otherwise indicating) a desired area in the dashboard template. Alternatively, a user can select the desired location by dragging the selected KPI (e.g., its identifier in the form of a KPI name), and dropping the selected KPI at the desired location in the dashboard template. For example, when the user selects the KPI, a default KPI widget is automatically displayed at a default location in the dashboard template. The user can then select the location by dragging and dropping the default KPI widget at the desired location. As will be discussed in greater detail below in conjunction with FIGS. 40-42 and FIGS. 44-46, a KPI widget is a KPI identifier that provides a numerical and/or visual representation of one or more values for the selected KPI. A KPI widget can be, for example, a Noel gauge, a spark line, a single value, a trend indicator, etc.

At block 3511, the computing machine receives a selection of one or more style settings for a KPI identifier (a KPI widget) to be displayed in the service-monitoring dashboard. For example, after the user selects the KPI, the user can provide input for creating and/or editing a title for the KPI. In one implementation, the computing machine causes the title that is already assigned to the selected KPI, for example via GUI 2200 in FIG. 22, to be displayed at the selected location in the dashboard template. In another example, after the user selects the KPI, the user is presented with available style settings, and the user can then select one or more of the style settings for the KPI widget to be displayed in the dashboard. In another example, in which a default KPI widget is displayed in response to the user selection of the KPI, the user can choose one or more of the available style setting(s) to replace or modify the default KPI widget. Style settings define how the KPI widget should be presented and can specify, for example, the shape of the widget, the size of the

widget, the name of the widget, the metric unit of a KPI value, and/or other visual characteristics of the widget. Some implementations for receiving a selection of style setting(s) for a KPI widget to be displayed in the dashboard are discussed in greater detail below in conjunction with FIG. 39. At block 3513, the computing machine causes display of a KPI identifier, such as a KPI widget, for the selected KPI at the selected location in the dashboard template. The KPI widget that is displayed in the dashboard template can be displayed using the selected style settings. The computing machine can receive further input (e.g., user input) for resizing a KPI widget via an input device (e.g., mouse, touch screen, etc.) For example, the computing device may receive user input via mouse device resizing (e.g., stretching, shrinking) the KPI widget.

FIG. 39 illustrates an example GUI 3900 facilitating user input for selecting a location in the dashboard template and style settings for a KPI widget, editing the service-monitoring dashboard by editing the dashboard template for the service-monitoring dashboard, and displaying the KPI widget in the dashboard template, in accordance with one or more implementations of the present disclosure. GUI 3900 includes a configuration interface 3906 to display a set of selectable thumbnail images (or icons or buttons) 3911 representing different types or styles of KPI widgets. The KPI widget styles can include, for example, and not limited to, a single value widget, a spark line widget, a Noel gauge widget, and a trend indicator widget. Configuration interface 3905 can display a single value widget thumbnail image 3907, a spark line widget thumbnail image 3908, a Noel gauge widget thumbnail image 3909, and a trend indicator widget thumbnail image 3910. For example, a user may have selected the Web Hosting service 3901, dragged the Web Hosting service 3901, and dropped the Web Hosting service 3901 on location 3905. The user may also have selected the CPU Usage KPI for the Web Hosting service 3901 and the Noel gauge widget thumbnail image 3909 to display the KPI widget for the CPU Usage KPI at the location 3905. In response, the computing machine can cause display of the Noel Gauge widget for the selected KPI (e.g., CPU Usage KPI) at the selected location (e.g., location 3905) in the dashboard template 3903. Some implementations of widgets for representing KPIs are discussed in greater detail below in conjunction with FIGS. 40-42 and FIGS. 44-46. In response to a user selection of a style setting for the KPI widget, one or more GUIs can be presented for customizing the selected KPI widget for the KPI. Input can be received via the GUIs to select a label for a KPI widget and the metric unit to be used for the KPI value with the KPI widget.

In one implementation, GUI 3900 includes an icon 3914 in the customization toolbar, which can be selected by a user, for defining one or more search queries. The search queries may produce results pertaining to one or more entities. For example, icon 3914 may be selected and an identifier 3918 for a search widget can be displayed in the dashboard template 3903. The identifier 3918 for the search widget can be the search widget itself, as illustrated in FIG. 39. The search widget can be a shape (e.g., box) and can display results (e.g., value produced by a corresponding search query) in the shape in the service-monitoring dashboard when the search query is executed for displaying the service-monitoring dashboard to a user.

The identifier 3918 can be displayed in a default location in the dashboard template 3903 and a user can optionally select a new location for the identifier 3918. The location of the identifier 3918 in the dashboard template specifies the location of the search widget in the service-monitoring dashboard when the service-monitoring dashboard is displayed to a user.

GUI **3900** can display a search definition box (e.g., box **3915**) that corresponds to the search query. A user can provide input for the criteria for the search query via the search definition box (e.g., box **3915**). For example, the search query may produce a stats count for a particular entity. The input pertaining to the search query is stored as part of the dashboard template. The search query can be executed when the service-monitoring dashboard is displayed to a user and the search widget can display the results from executing the search query.

Referring to FIG. **35**, in one implementation, the computing machine receives input (e.g., user input), via the dashboard-creation graphical interface, of a time range to use for the KPI widget, editing the service-monitoring dashboard, and clearing data in the dashboard template.

At block **3515**, the computing machine stores the resulting dashboard template in a data store. The dashboard template can be saved in response to a user request. For example, a request to save the dashboard template may be received upon selection of a save button (e.g., save button **3612** in GUI **3600** of FIG. **36**).

Referring to FIG. **35**, at block **3517**, the computing machine can receive a user request for a service-monitoring dashboard, and can then generate and cause display of the service-monitoring dashboard based on the dashboard template at block **3519**. Some implementations for causing display of a service-monitoring dashboard based on the dashboard template are discussed in greater detail below in conjunction with FIG. **47**.

FIG. **40** illustrates an example Noel gauge widget **4000**, in accordance with one or more implementations of the present disclosure. Noel gauge widget **4000** can have a shape **4001** with an empty space **4002** and with one end **4004** corresponding to a minimum KPI value and the other end **4006** corresponding to a maximum KPI value. The minimum value and maximum value can be user-defined values, for example, received via fields **3116, 3120** in GUI **3100** in FIG. **31A**, as discussed above. Referring to FIG. **40**, the value produced by the search query defining the KPI can be represented by filling in the empty space **4002** of the shape **4001**. This filler can be displayed using a color **4003** to represent the current state (e.g., normal, warning, critical) of the KPI according to the value produced by the search query. The color can be based on input received when one or more thresholds were created for the KPI. The Noel gauge widget **4000** can also display the actual value **4007** produced by the search query defining the KPI. The value **4007** can be of a nominal color or can be of a color representative of the state to which the value produced by the search query corresponds. A user can provide input, via the dashboard-creation graphical interface, indicating whether to apply a nominal color or color representative of the state.

The Noel gauge widget **4000** can display a label **4005** (e.g., Request Response Time) to describe the KPI and the metric unit **4009** (e.g., ms (milliseconds)) used for the KPI value. If the KPI value **4007** exceeds the maximum value represented by the second end **4006** of the shape **4001** of the Noel gauge widget **4000**, the shape **4001** is displayed as being fully filled and can include an additional visual indicator representing that the KPI value **4007** exceeded the maximum value represented by the second end **4006** of the shape **4001** of the Noel gauge widget **4000**.

The value **4007** can be produced by executing the search query of the KPI. The execution can be real-time (continuous execution until interrupted) or relative (based on a specific request or scheduled time). In addition, the machine data used by the search query to produce each value can be based on a

time range. The time range can be user-defined time range. For example, before displaying a service-monitoring dashboard generated based on the dashboard template, a user can provide input specifying the time range. The input can be received, for example, via a drop-down menu **3912** in GUI **3900** in FIG. **39**. The initial time range, received via GUI **3900**, can be stored with the dashboard template in a data store and subsequently used for producing the values for the KPI to be displayed in the service-monitoring dashboard.

When drop-down menu **3912** is selected by a user, GUI **4300** in FIG. **43** can be displayed. FIG. **43** illustrates an example GUI **4300** for facilitating user input specifying a time range to use when executing a search query defining a KPI, in accordance with one or more implementations of the present disclosure. For real-time execution, for example, used to update the service-monitoring dashboard in real-time, the time range for machine data can be a specified time window (e.g., 30-second window, 1-minute window, 1-hour window, etc.) from the execution time (e.g., each time the query is executed, the events with timestamps within the specified time window from the query execution time will be used). For relative execution, the time range can be historical (e.g., yesterday, previous week, etc.) or based on a specified time window from the requested time or scheduled time (e.g., last 15 minutes, last 4 hours, etc.). For example, the historical time range "Yesterday" **4304** can be selected for relative execution. In another example, the window time range "Last 15 minutes" **4305** can be selected for relative execution.

Referring to FIG. **40**, the KPI may be for Request Response Time for a Web Hosting service. The time range "Last 15 minutes" may be selected for the service-monitoring dashboard presented to a user, and the value **4007** (e.g., 1.41) produced by the search query defining the Request Response Time KPI can be the average response time using the last 15 minutes of machine data associated with the entities providing the Web Hosting service from the time of the request. FIG. **42** illustrates an example GUI **4200** illustrating a search query and a search result for a Noel gauge widget, a single value widget, and a trend indicator widget, in accordance with one or more implementations of the present disclosure. A single value widget is discussed in greater detail below in conjunction with FIG. **41**. A trend indicator widget is discussed in greater detail below in conjunction with FIG. **46**. Referring to FIG. **42**, the KPI may be for Request Response Time. The KPI may be defined by a search query **4501** that outputs a search result having a single value **4203** (e.g., 1.41) for a Noel gauge widget, a single value widget, and/or a trend indicator widget. The search query **4201** can include a statistical function **4205** (e.g., average) to produce the single value (e.g., value **4203**) to represent response time using machine data from the Last 15 minutes **4207**.

FIG. **41** illustrates an example single value widget **4100**, in accordance with one or more implementations of the present disclosure. Single value widget **4100** can include the value **4107**, produced by the search query defining the KPI, in a shape **4101** (e.g., box). The shape can be colored using a color **4103** representative of the state (e.g., normal, warning, critical) to which the value produced by the search query corresponds. The value **4107** can be also colored using a nominal color or a color representative of the state to which the value produced by the search query corresponds. The single value widget **4100** can display a label to describe the KPI and the metric unit used for the KPI. A user can provide input, via the dashboard-creation graphical interface, indicating whether to apply a nominal color or color representative of the state.

The machine data used by the search query to produce the value **4107** is based on a time range (e.g., user selected time

range). For example, the KPI may be for Request Response Time for a Web Hosting service. The time range "Last 15 minutes" may be selected for the service-monitoring dashboard presented to a user. The value **4107** (e.g., **1.41**) produced by the search query defining the Request Response Time KPI can be the average response time using the last 15 minutes of machine data associated with the entities providing the Web Hosting service from the time of the request.

FIG. **44** illustrates spark line widget **4400**, in accordance with one or more implementations of the present disclosure. Spark line widget **4400** can include two shapes (e.g., box **4405** and rectangular box **4402**). One shape (e.g., box **4405**) of the spark line widget **4400** can include a value **4407**, which is described in greater detail below. The shape (e.g., box **4405**) can be colored using a color **4406** representative of the state (e.g., normal, warning, critical) to which the value **4407** corresponds. The value **4407** can be also be colored using a nominal color or a color representative of the state to which the value **4407** corresponds. A user can provide input, via the dashboard-creation graphical interface, indicating whether to apply a nominal color or color representative of the state.

Another shape (e.g., rectangular box **4402**) in the spark line widget **4400** can include a graph **4401** (e.g., line graph), which is described in greater detail below, that includes multiple data points. The shape (e.g., rectangular box **4402**) containing the graph **4401** can be colored using a color representative of the state (e.g., normal, warning, critical) of which a corresponding data point (e.g., latest data point) falls into. The graph **4401** can be colored using a color representative of the state (e.g., normal, warning, critical) of which a corresponding data point falls into. For example, the graph **4401** may be a line graph that transitions between green, yellow, red, depending on the value of a data point in the line graph. In one implementation, input (e.g., user input) can be received, via the service-monitoring dashboard, of a selection device hovering over a particular point in the line graph, and information (e.g., data value, time, and color) corresponding to the particular point in the line graph can be displayed in the service-monitoring dashboard, for example, in the spark line widget **4400**. The spark line widget **4400** can display a label to describe the KPI and the metric unit used for the KPI.

The spark line widget **4400** is showing data in a time series graph with the graph **4401**, as compared to a single value widget (e.g., single value widget **4100**) and a Noel gauge widget (e.g., Noel gauge widget **4000**) that display a single data point, for example as illustrated in FIG. **42**. The data points in the graph **4401** can represent what the values, produced by the search query defining the KPI, have been over a time range (e.g., time range selected in GUI **4300**). FIG. **45** illustrates an example GUI **4500** illustrating a search query and search results for a spark line widget, in accordance with one or more implementations of the present disclosure. The KPI may be for Request Response Time. The KPI may be defined by a search query **4501** that produces multiple values, for example, to be used for a spark line widget. A user may have selected a time range of "Last 15 minutes" **4507** (e.g., time range selected in GUI **4300**). The machine data used by the search query **4501** to produce the search results can be based on the last 15 minutes. For example, the search results can include a value for each minute in the last 15 minutes. The values **4503** in the search results can be used as data points to plot a graph (e.g., graph **4401** in FIG. **44**) in the spark line widget. Referring to FIG. **44**, the graph **4401** is from data over a period of time (e.g., Last 15 minutes). The graph **4401** is made of data points (e.g., 15 values **4503** in search results in FIG. **45**). Each data point is an aggregate from the data for a shorter period of time (e.g., unit of time). For example, if the

time range "Last 15 minutes" is selected, each data point in the graph **4401** represents a unit of time in the last 15 minutes. For example, the unit of time may be one minute, and the graph contains 15 data points, one for each minute for the last 15 minutes. Each data point can be the average response time (e.g., avg(spent) in search query **4501** in FIG. **45**) for the corresponding minute. In another example, if the time range "Last 4 hours" is selected, and the unit of time used for the graph **4401** is 15 minutes, then the graph **4401** would be made from 16 data points.

In one implementation, the value **4407** in the other shape (e.g., box **4405**) in the spark line widget **4400** represents the latest value in the time range. For example, the value **4407** (e.g., **1.32**) can represent the last data point **4403** in the graph **4401**. If the time range "Last 15 minutes" is selected, the value **4407** (e.g., **1.32**) can represent the average response time of the data in that last minute of the 15 minute time range.

In another implementation, the value **4407** is the first data point in the graph **4401**. In another implementation, the value **4407** represents an aggregate of the data in the graph **4401**. For example, a statistical function can be performed on using the data points for the time range (e.g., Last 15 minutes) for the value **4407**. For example, the value **4407** may be the average of all of the points in the graph **4401**, the maximum value from all of the points in the graph **4401**, the mean of all of the points in the graph **4401**. Input (e.g., user input) can be received, for example, via the dashboard-creation graphical interface, specifying type (e.g. latest, first, average, maximum, mean) of value to be represented by value **4407**.

FIG. **46** illustrates a trend indicator widget **4600**, in accordance with one or more implementations of the present disclosure. Trend indicator widget **4600** can include a shape **4601** (e.g., rectangular box) that includes a value **4607**, produced by the search query defining the KPI, in another shape **4601** (e.g., box) and an arrow **4605**. The shape **4601** containing the value **4607** can be colored using a color **4603** representative of the state (e.g., normal, warning, critical) of which the value **4607** produced by the search query falls into. The value **4607** can be of a nominal color or can be of a color representative of the state for which the value produced by the search query falls into. A user can provide input, via the dashboard-creation graphical interface, indicating whether to apply a nominal color or color representative of the state. The trend indicator widget **4600** can display a label to describe the KPI and the metric unit used for the KPI.

The arrow **4605** can indicate a trend pertaining to the KPI by pointing in a direction. For example, the arrow **4605** can point in a general up direction to indicate a positive or increasing trend, the arrow **4605** can point in a general down direction to indicate a negative or decreasing trend, or the arrow **4605** can point in a general horizontal direction to indicate no change in the KPI. The direction of the arrow **4605** in the trend indicator widget **4600** may change when a KPI is being updated, for example, in a service-monitoring dashboard, depending on the current trend at the time the KPI is being updated.

In one implementation, a color is assigned to each trend (e.g., increasing trend, decreasing trend). The arrow **4605** can be of a nominal color or can be of a color representative of the determined trend. A user can provide input, via the dashboard-creation graphical interface, indicating whether to apply a nominal color or color representative of the trend. The shape **4607** can be of a nominal color or can be of a color representative of the determined trend. A user can provide input, via the dashboard-creation graphical interface, indicating whether to apply a nominal color or color representative of the trend.

In one implementation, the trend represented by the arrow **4605** is of whether the value **4607** has been increasing or decreasing in a selected time range relative to the last time the KPI was calculated. For example, if the time range “Last 15 minutes” is selected, the average of the data points of the last 15 minutes is calculated, and the arrow **4605** can indicate whether the average of the data points of the last 15 minutes is greater than the average calculated from the time range (e.g., 15 minutes) prior. In one implementation, the trend indicator widget **4600** includes a percentage indicator indicating a percentage of the value **4607** increasing or decreasing in a selected time range relative to the last time the KPI was calculated.

In another implementation, the arrow **4605** indicates whether the last value for the last data point in the last 15 minutes is greater than the value immediately before the last data point.

The machine data used by the search query to produce the value **4607** is based on a time range (e.g., user selected time range). For example, the KPI may be for Request Response Time for a Web Hosting service. The time range “Last 15 minutes” may be selected for the service-monitoring dashboard presented to a user. The value **4607** (e.g., **1.41**) produced by the search query defining the Request Response Time KPI can be the average response time using the last 15 minutes of machine data associated with the entities providing the Web Hosting service from the time of the request.

As discussed above, once the dashboard template is defined, it can be saved, and then used to generate a service-monitoring dashboard for display. The dashboard template can identify the KPIs selected for the service-monitoring dashboard, KPI widgets to be displayed for the KPIs in the service-monitoring dashboard, locations in the service-monitoring dashboard for displaying the KPI widgets, visual characteristics of the KPI widgets, and other information (e.g., the background image for the service-monitoring dashboard, an initial time range for the service-monitoring dashboard).

FIG. 47A is a flow diagram of an implementation of a method **4750** for creating and causing for display a service-monitoring dashboard, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method is performed by the client computing machine. In another implementation, the method is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block **4751**, the computing machine identifies one or more key performance indicators (KPIs) for one or more services to be monitored via a service-monitoring dashboard. A service can be provided by one or more entities. Each entity can be associated with machine data. The machine data can include unstructured data, log data, and/or wire data. The machine data associated with an entity can include data collected from an API for software that monitors that entity.

A KPI can be defined by a search query that derives one or more values from machine data associated with the one or more entities that provide the service. Each KPI can be defined by a search query that is either entered by a user or generated through a graphical interface. In one implementation, the computing machine accesses a dashboard template that is stored in a data store that includes information identifying the KPIs to be displayed in the service-monitoring dashboard. In one implementation, the dashboard template specifies a service definition that associates the service with

the entities providing the service, specifies the KPIs of the service, and also specifies the search queries for the KPIs. As discussed above, the search query defining a KPI can derive one or more values for the KPI using a late-binding schema that it applies to machine data. In some implementations, the service definition identified by the dashboard template can also include information on predefined states for a KPI and various visual indicators that should be used to illustrate states of the KPI in the dashboard.

The computing machine can optionally receive input (e.g., user input) identifying one or more ad hoc searches to be monitored via the service-monitoring dashboard without selecting services or KPIs.

At block **4753**, the computing machine identifies a time range. The time range can be a default time range or a time range specified in the dashboard template. The machine data can be represented as events. The time range can be used to indicate which events to use for the search queries for the identified KPIs.

At block **4755**, for each KPI, the computing machine identifies a KPI widget style to represent the respective KPI. In one implementation, the computing machine accesses the dashboard template that includes information identifying the KPI widget style to use for each KPI. As discussed above, examples of KPI widget styles can include a Noel gauge widget style, a single value widget style, a spark line widget style, and a trend indicator widget style. The computing machine can also obtain, from the dashboard template, additional visual characteristics for each KPI widget, such as, the name of the widget, the metric unit of the KPI value, settings for using nominal colors or colors to represent states and/or trends, the type of value to be represented in KPI widget (e.g., the type of value to be represented by value **4407** in a spark line widget), etc.

The KPIs widget styles can display data differently for representing a respective KPI. The computing machine can produce a set of values and/or a value, depending on the KPI widget style for a respective KPI. If the KPI widget style represents the respective KPI using values for multiple points in time in the time range, method **4750** proceeds to block **4757**. Alternatively, if the KPI widget style represents the respective KPI using a single value during the time range, method **4750** proceeds to block **4759**.

At block **4759**, if the KPI widget style represents the respective KPI using a single value, the computing machine causes a value to be produced from a set of machine data or events whose timestamps are within the time range. The value may be a statistic calculated based on one or more values extracted from a specific field in the set of machine data or events when the search query is executed. The statistic may be an average of the extracted values, a mean of the extracted values, a maximum of the extracted values, a last value of the extracted values, etc. A single value widget style, a Noel gauge widget style, and trend indicator widget style can represent a KPI using a single value.

The search query that defines a respective KPI may produce a single value which a KPI widget style can use. The computing machine can cause the search query to be executed to produce the value. The computing machine can send the search query to an event processing system. As discussed above, machine data can be represented as events. The machine data used to derive the one or more KPI values can be identifiable on a per entity basis by referencing entity definitions that are aggregated into a service definition corresponding to the service whose performance is reflected by the KPI.

The event processing system can access events with time stamps falling within the time period specified by the time

range, identify which of those events should be used (e.g., from the one or more entity definitions in the service definition for the service whose performance is reflected by the KPI), produce the result (e.g., single value) using the identified events, and send the result to the computing machine. The computing machine can receive the result and store the result in a data store.

At block 4757, if the KPI widget style represents the respective KPI using a set of values, the computing machine causes a set of values for multiple points in time in the time range to be produced. A spark line widget style can represent a KPI using a set of values. Each value in the set of values can represent an aggregate of data in a unit of time in the time range. For example, if the time range is "Last 15 minutes", and the unit of time is one minute, then each value in the set of values is an aggregate of the data in one minute in the last 15 minutes.

If the search query that defines a respective KPI produces a single value instead of a set of multiple values as required by the KPI widget style (e.g., for the graph of the spark line widget), the computing machine can modify the search query to produce the set of values (e.g., for the graph of the spark line widget). The computing machine can cause the search query or modified search query to be executed to produce the set of values. The computing machine can send the search query or modified search query to an event processing system. The event processing system can access events with time stamps falling within the time period specified by the time range, identify which of those events should be used, produce the results (e.g., set of values) using the identified events, and send the results to the computing machine. The computing machine can store the results in a data store.

At block 4761, for each KPI, the computing machine generates the KPI widget using the KPI widget style and the value or set of values produced for the respective KPI. For example, if a KPI is being represented by a spark line widget style, the computing machine generates the spark line widget using a set of values produced for the KPI to populate the graph in the spark line widget. The computing machine also generates the value (e.g., value 4407 in spark line widget 4400 in FIG. 44) for the spark line widget based on the dashboard template. The dashboard template can store the selection of the type of value that is to be represented in a spark line widget. For example, the value may represent the first data point in the graph, the last data point the graph, an average of all of the points in the graph, the maximum value from all of the points in the graph, or the mean of all of the points in the graph.

In addition, in some implementations, the computing machine can obtain KPI state information (e.g., from the service definition) specifying KPI states, a range of values for each state, and a visual characteristic (e.g., color) associated with each state. The computing machine can then determine the current state of each KPI using the value or set of values produced for the respective KPI, and the state information of the respective KPI. Based on the current state of the KPI, a specific visual characteristic (e.g., color) can be used for displaying the KPI widget of the KPI, as discussed in more detail above.

At block 4763, the computing machine generates a service-monitoring dashboard with the KPI widgets for the KPIs using the dashboard template and the KPI values produced by the respective search queries. In one implementation, the computing machine accesses a dashboard template that is stored in a data store that includes information identifying the KPIs to be displayed in the service-monitoring dashboard. As discussed above, the dashboard template defines locations for

placing the KPI widgets, and can also specify a background image over which the KPI widgets can be placed.

At block 4765, the computing machine causes display of the service-monitoring dashboard with the KPI widgets for the KPIs. Each KPI widget provides a numerical and/or graphical representation of one or more values for a corresponding KPI. Each KPI widget indicates how an aspect of the service is performing at one or more points in time. For example, each KPI widget can display a current KPI value, and indicate the current state of the KPI using a visual characteristic associated with the current state. In one implementation, the service-monitoring dashboard is displayed in a viewing/investigation mode based on a user selection to view the service-monitoring dashboard is displayed in the viewing/investigation mode.

At block 4767, the computing machine optionally receives a request for detailed information for one or more KPIs in the service-monitoring dashboard. The request can be received, for example, from a selection (e.g., user selection) of one or more KPI widgets in the service-monitoring dashboard.

At block 4759, the computing machine causes display of the detailed information for the one or more KPIs. In one implementation, the computing machine causes the display of a deep dive visual interface, which includes detailed information for the one or more KPIs. A deep dive visual interface is described in greater detail below in conjunction with FIG. 50A.

The service-monitoring dashboard may allow a user to change a time range. In response, the computing machine can send the search query and the new time range to an event processing system. The event processing system can access events with time stamps falling within the time period specified by the new time range, identify which of those events should be used, produce the result (e.g., one or more values) using the identified events, and send the result to the computing machine. The computing machine can then cause the service-monitoring dashboard to be updated with new values and modified visual representations of the KPI widgets.

FIG. 47B illustrates an example service-monitoring dashboard GUI 4700 that is displayed based on the dashboard template, in accordance with one or more implementations of the present disclosure. GUI 4700 includes a user selected background image 4702 and one or more KPI widgets for one or more services that are displayed over the background image 4702. GUI 4700 can include other elements, such as, and not limited to, text, boxes, connections, and widgets for ad hoc searches. Each KPI widget provides a numerical or graphical representation of one or more values for a corresponding key performance indicator (KPI) indicating how an aspect of a respective service is performing at one or more points in time. For example, GUI 4700 includes a spark line widget 4718 which may be for an aspect of Service-B, and a Noel gauge widget 4708 which may be for another aspect of Service-B. In some implementations, the appearance of the widgets 4718 and 4708 (as well as other widgets in the GUI 4700) can reflect the current state of the respective KPI (e.g., based on color or other visual characteristic).

Each service is provided by one or more entities. Each entity is associated with machine data. The machine data can include for example, and is not limited to, unstructured data, log data, and wire data. The machine data that is associated with an entity can include data collected from an API for software that monitors that entity. The machine data used to derive the one or more values represented by a KPI is identifiable on a per entity basis by referencing entity definitions that are aggregated into a service definition corresponding to the service whose performance is reflected by the KPI.

57

Each KPI is defined by a search query that derives the one or more values represented by the corresponding KPI widget from the machine data associated with the one or more entities that provide the service whose performance is reflected by the KPI. The search query for a KPI can derive the one or more values for the KPI it defines using a late-binding schema that it applies to machine data.

In one implementation, the GUI **4700** includes one or more search result widgets (e.g., widget **4712**) displaying a value produced by a respective search query, as specified by the dashboard template. For example, widget **4712** may represent the results of a search query producing a stats count for a particular entity.

In one implementation, GUI **4700** facilitates user input for displaying detailed information for one or more KPIs. A user can select one or more KPI widgets to request detailed information for the KPIs represented by the selected KPI widgets. The detailed information for each selected KPI can include values for points in time during the period of time. The detailed information can be displayed via one or more deep dive visual interfaces. A deep dive visual interface is described in greater detail below in conjunction with FIG. **50A**.

Referring to FIG. **47**, GUI **4700** facilitates user input for changing a time range. The machine data used by a search query to produce a value for a KPI is based on a time range. As described above in conjunction with FIG. **43**, the time range can be a user-defined time range. For example, if the time range “Last 15 minutes” is selected, the last 15 minutes would be an aggregation period for producing the value. GUI **4700** can be updated with one or more KPI widgets that each represent one or more values for a corresponding KPI indicating how a service provided is performing at one or more points in time based on the change to the time range.

FIG. **48** describes an example home page GUI **4800** for service-level monitoring, in accordance with one or more implementations of the present disclosure. GUI **4800** can include one or more tiles each representing a service-monitoring dashboard. The GUI **4800** can also include one or more tiles representing a service-related alarm, or the value of a particular KPI. In one implementation, a tile is a thumbnail image of a service-monitoring dashboard. When a service-monitoring dashboard is created, a tile representing the service-monitoring dashboard can be displayed in the GUI **4800**. GUI **4800** can facilitate user input for selecting to view a service-monitoring dashboard. For example, a user may select a dashboard tile **4805**, and GUI **4700** in FIG. **47** can be displayed in response to the selection. GUI **4800** can include tiles representing the most recently accessed dashboards, and user selected favorites of dashboards.

FIG. **49** describes an example home page GUI **4900** for service-level monitoring, in accordance with one or more implementations of the present disclosure. GUI **4900** can include one or more tiles representing a deep dive. In one implementation, a tile is a thumbnail image of a deep dive. When a deep dive is created, a tile representing the deep dive can be displayed in the GUI **4900**. GUI **4900** can facilitate user input for selecting to view a deep dive. For example, a user may select a deep dive tile **4907**, and the visual interface **5300** in FIG. **55** can be displayed in response to the selection. GUI **4900** can include tiles representing the most recently accessed deep dives, and user selected favorites of deep dives. Example Deep Dive

Implementations of the present disclosure provide a GUI that provides in-depth information about multiple KPIs of the same service or different services. This GUI referred to herein as a deep dive displays time-based graphical visualizations

58

corresponding to the multiple KPIs to allow a user to visually correlate the KPIs over a defined period of time.

FIG. **50A** is a flow diagram of an implementation of a method for creating a visual interface displaying graphical visualizations of KPI values along time-based graph lanes, in accordance with one or more implementations of the present disclosure. The method may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one implementation, the method **5000** is performed by a client computing machine. In another implementation, the method **5000** is performed by a server computing machine coupled to the client computing machine over one or more networks.

At block **5001**, the computing machine receives a selection of KPIs that each indicates a different aspect of how a service (e.g., a web hosting service, an email service, a database service) provided by one or more entities (e.g., host machines, virtual machines, switches, firewalls, routers, sensors, etc.) is performing. As discussed above, each of these entities produces machine data or has its operation reflected in machine data not produced by that entity (e.g., machine data collected from an API for software that monitors that entity while running on another entity). Each KPI is defined by a different search query that derives one or more values from the machine data pertaining to the entities providing the service. Each of the derived values is associated with a point in time and represents the aspect of how the service is performing at the associated point in time. In one implementation, the KPIs are selected by a user using GUIs described below in connection with FIGS. **51**, **52** and **57-63**.

At block **5003**, the computing machine derives the value(s) for each of the selected KPIs from the machine data pertaining to the entities providing the service. In one implementation, the computing machine executes a search query of a respective KPI to derive the value(s) for that KPI from the machine data.

At block **5005**, the computing machine causes display of a graphical visualization of the derived KPI values along a time-based graph lane for each of the selected KPIs. In one implementation, the graph lanes for the selected KPIs are parallel to each other and the graphical visualizations in the graph lanes are all calibrated to the same time scale. In one implementation, the graphical visualizations are displayed in the visual interfaces described below in connection with FIGS. **53-56** and **64A-70**.

FIG. **50B** is a flow diagram of an implementation of a method for generating a graphical visualization of KPI values along a time-based graph lane, in accordance with one or more implementations of the present disclosure.

At block **5011**, the computing machine receives a request to create a graph for a KPI. Depending on the implementation, the request can be made by a user from service-monitoring dashboard GUI **4700** or from a GUI **5100** for creating a visual interface, as described below with respect to FIG. **51**. At block **5013**, the computing machine displays the available services that are being monitored, and at block **5015**, receives a selection of one of the available services. At block **5017**, the computing machine displays the KPIs associated with the selected service, and at block **5019**, receives a selection of one of the associated KPIs. In one implementation, the KPIs are selected by a user using GUIs described below in connection with FIGS. **51**, **52** and **57-63**. At block **5021**, the computing machine uses a service definition of the selected service to identify a search query corresponding to the selected KPI. At block **5023**, the computing machine determines if there are

59

more KPI graphs to create. If the user desires to create additional graphs, the method returns to block 5013 and repeats the operations of blocks 5013-5021 for each additional graph.

If there are no more KPI graphs to create, at block 5025, the computing machine identifies a time range. The time range can be defined based on a user input, which can include, e.g., identification of a relative time or absolute time, perhaps entered through user interface controls. The time range can include a portion (or all) of a time period, where the time period corresponds to one used to indicate which values of the KPI to retrieve from a data store. In one implementation, the time range is selected by a user using GUIs described below in connection with FIGS. 54 and 63. At block 5027, the computing device creates a time axis reflecting the identified time range. The time axis may run parallel to at least one graph lane in the create visual interface and may include an indication of the amount of time represented by a time scale for the visual interface (e.g., "Viewport: 1 h 1 m" indicating that the graphical visualizations in the graph lanes display KPI values for a time range of one hour and one minute).

At block 5029, the computing device executes the search query corresponding to each KPI and stores the resulting KPI dataset values for the selected time range. At block 5031, the computing device determines the maximum and minimum values of the KPI for the selected time range and at block 5033 creates a graph lane in the visual interface for each KPI using the maximum and minimum values as the height of the lane. In one implementation, a vertical scale for each lane may be automatically selected using the maximum and minimum KPI values during the current time range, such that the maximum value appears at or near the top of the lane and the minimum value appears at or near the bottom of the lane. The intermediate values between the maximum and minimum may be scaled accordingly.

At block 5035, the computing device creates a graphical visualization for each lane using the KPI values during the selected time period and selected visual characteristics. In one implementation, the KPI values are plotted over the time range in a time-based graph lane. The graphical visualization may be generated according to an identified graph type and graph color, as well as any other identified visual characteristics. At block 5037, the computing device calibrates the graphical visualizations to a same time scale, such that the graphical visualization in each lane of the visual interface represents KPI data over the same period of time.

Blocks 5025-5037 can be repeated for a new time range. Such repetition can occur, e.g., after detecting an input corresponding to an identification of a new time range. The generation of a new graphical visualization can include modification of an existing graphical visualization.

FIG. 51 illustrates an example GUI 5100 for creating a visual interface displaying graphical visualizations of KPI values along time-based graph lanes, in accordance with one or more implementations of the present disclosure. The GUI 5100 can receive user input for a number of input fields 5102, 5104 and selection of selection buttons 5106. For example, input field 5102 can receive a title for the visual interface being created. Input field 5104 may receive a description of the visual interface. The input to input fields 5102 and 5104 may be optional in one implementation, such that it is not absolutely required for creation of the visual interface. Input to fields 5102 and 5104 may be helpful, however, in identifying the visual interface once it is created. In one implementation, if a title is not received in input fields 5102 and 5104, the computing machine may assign a default title to the created visual interface. Selection buttons 5106 may receive input pertaining to an access permission for the visual inter-

60

face being created. In one implementation, the user may select an access permission of either "Private," indicating that the visual interface being created will not be accessible to any other users of the system instead being reserved for private use by the user, or "Shared," indicating that once created, the visual interface will be accessible to other users of the system. Upon, the optional entering of title and description into fields 5102 and 5104 and the selection of an access permission using buttons 5106, the selection of button 5108 may initiate creation of the visual interface. In one implementation, in addition to "Private" or "Shared" there may be additional or intermediate levels of access permissions. For example, certain individuals or groups of individuals may be granted access or denied access to a given visual interface. There may be a role based access control system where individuals assigned to a certain role are granted access or denied access.

FIG. 52 illustrates an example GUI 5200 for adding a graphical visualization of KPI values along a time-based graph lane to a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, in response to the creation of a visual interface using GUI 5100, the system presents GUI 5200 in order to add graphical visualizations to the visual interface. The graphical visualizations correspond to KPIs and are displayed along a time-based graph lane in the visual interface.

In one example, GUI 5200 can receive user input for a number of input fields 5202, 5204, 5212, selections from drop down menus 5206, 5208, and/or selection of selection buttons 5210 or link 5214. For example, input field 5202 can receive a title for the graphical visualization being added. Input field 5204 may receive a subtitle or description of the graphical visualization. The input to input fields 5202 and 5204 may be optional in one implementation, such that it is not absolutely required for addition of the graphical visualization. Input to fields 5202 and 5204 may be helpful, however, in identifying the graphical visualization once it is added to the visual interface. In one implementation, if a title is not received in input fields 5202 or 5204, the computing machine may assign a default title to the graphical visualization being added.

Drop down menu 5206 can be used to receive a selection of a graph style, and drop down menu 5208 can be used to receive a selection of a graph color for the graphical visualization being added. Additional details with respect to selection of the graph style and the graph color for the graphical visualization are described below in connection with FIGS. 57 and 58.

Selection buttons 5210 may receive input pertaining to a search source for the graphical visualization being added. In one implementation, the user may select search source of "Ad Hoc," "Data Model" or "KPI." Additional details with respect to selection of the search source for the graphical visualization are described below in connection with FIGS. 57, 59 and 60. Input field 5212 may receive a user-input search query or display a search query associated with the selected search source 5210. Selection of link 5214 may indicate that the user wants to execute the search query in input field 5212. When a search query is executed, the search query can produce one or more values that satisfy the search criteria for the search query. Upon the entering of data and the selection menu items, the selection of button 5216 may initiate the addition of the graphical visualization to the visual interface.

FIG. 53 illustrates an example of a visual interface 5300 with time-based graph lanes for displaying graphical visualizations, in accordance with one or more implementations of the present disclosure. In one example, the visual interface 5300 includes three time-based graph lanes 5302, 5304, 5306. These graph lanes may correspond to the graphical visualiza-



61

tions of KPI values added to the visual interface using GUI 5200 as described above. Each of the graph lanes 5302, 5304, 5306 can display a graphical visualization for corresponding KPI values over a time range. Initially the lanes 5302, 5304, 5306 may not include the graphical visualizations until a time range is selected using drop down menu 5308. Additional details with respect to selection of the time range from drop down menu 5308 are described below in connection with FIG. 63. In another implementation, a default time range may be automatically selected and the graphical visualizations may be displayed in lanes 5302, 5304, 5306.

FIG. 54 illustrates an example of a visual interface 5300 displaying graphical visualizations of KPI values along time-based graph lanes, in accordance with one or more implementations of the present disclosure. In one implementation, each of the time-based graph lanes 5302, 5304, 5306 include a visual representation of corresponding KPI values. The visual representations in each lane may be of different graph styles and/or colors or have the same graph styles and/or colors. For example, lane 5302 includes a bar chart, lane 5304 includes a line graph and lane 5306 includes a bar chart. The graph type and graph color of the visual representation in each lane may be selected using GUI 5200, as described above. Depending on the implementation, the KPIs represented by the graphical visualizations may correspond to different services or may correspond to the same service. In one implementation, multiple of the KPIs may correspond to the same service, while one or more other KPIs may correspond to a different service.

The graphical visualizations in each lane 5302, 5304, 5306 can all be calibrated to the same time scale. That is, each graphical visualization corresponds to a different KPI reflecting how a service is performing over a given time range. The time range can be reflected by a time axis 5410 for the graphical visualizations that runs parallel to at least one graph lane. The time axis 5410 may include an indication of the amount of time represented by the time scale (e.g., “Viewport: 1 h 1 m” indicating that the graphical visualizations in graph lanes 5302, 5304, 5306 display KPI values for a time range of one hour and one minute), and an indication of the actual time of day represented by the time scale (e.g., “12:30, 12:45, 01 PM, 01:15”). In one implementation, a bar running parallel to the time lanes including the indication of the amount of time represented by the time scale (e.g., “Viewport: 1 h 1 m”) is highlighted for an entire length of time axis 5410 to indicate that the current portion of the time range being viewed includes the entire time range. In other implementations, when only a subset of the time range is being viewed, the bar may be highlighted for a proportional subset of the length of time axis 5410 and only in a location along time axis 5410 corresponding to the subset. In one implementation, at least a portion of the time axis 5410 is displayed both above and below the graph lanes 5302, 5304, 5306. In one implementation, an indicator associated with drop down menu 5308 also indicates the selected time range (e.g., “Last 60 minutes”) for the graphical visualizations.

In one implementation, when one of graph lanes 5302, 5304, 5306 is selected (e.g., by hovering the cursor over the lane), a grab handle 5412 is displayed in association with the selected lane 5302. When user interaction with grab handle 5412 is detected (e.g., by click and hold of a mouse button), the graph lanes may be re-ordered in visual interface 5300. For example, a user may use grab handle 5412 to move lane 5302 to a different location in visual interface 5300 with respect to the other lanes 5304, 5306, such as between lanes 5304 and 5306 or below lanes 5304 and 5306. When another lane is selected, a corresponding grab handle may be displayed for the selected lane and used to detect an interaction

62

of a user indicative of an instruction to re-order the graph lanes. In one implementation, a grab handle 5412 is only displayed when the corresponding lane 5302 is selected, and hidden from view when the lane is not selected.

While the horizontal axis of each lane is scaled according to the selected time range, and may be the same for each of the lanes 5302, 5304, 5306, a scale for the vertical axis of each lane may be determined individually. In one implementation, a scale for the vertical axis of each lane may be automatically selected such that the graphical visualization spans most or all of a width/height of the lane. In one implementation, the scale may be determined using the maximum and minimum values reflected by the graphical visualization for the corresponding KPI during the current time range, such that the maximum value appears at or near the top of the lane and the minimum value appears at or near the bottom of the lane. The intermediate values between the maximum and minimum may be scaled accordingly. In one implementation, a search query associated with the KPI is executed for a selected period of time. The results of the query return a dataset of KPI values, as shown in FIG. 45. The maximum and minimum values from this dataset can be determined and used to scale the graphical visualization so that most or all of the lane is utilized to display the graphical visualization.

FIG. 55A illustrates an example of a visual interface 5300 with a user manipulable visual indicator 5514 spanning across the time-based graph lanes, in accordance with one or more implementations of the present disclosure. Visual indicator 5514, also referred to herein as a “lane inspector,” may include, for example, a line or other indicator that spans vertically across the graph lanes 5302, 5304, 5306 at a given point in time along time axis 5410. The visual indicator 5514 may be user manipulable such that it may be moved along time axis 5410 to different points. For example, visual indicator 5514 may slide back and forth along the lengths of graph lanes 5302, 5304, 5306 and time axis 5410 in response to user input received with a mouse, touchpad, touchscreen, etc.

In one implementation, visual indicator 5514 includes a display of the point in time at which it is currently located. In the illustrated example, the time associated with visual indicator 5514 is “12:44:43 PM.” In one implementation, visual indicator 5514 further includes a display of a value reflected in each of the graphical visualizations for the different KPIs at the current point in time illustrated by visual indicator 5514. In the illustrated example, the value of the graphical visualization in lane 5302 is “3.65,” the value of the graphical visualization in lane 5304 is “60,” and the value of the graphical visualization in lane 5306 is “0.” In one implementation, units for the values of the KPIs are not displayed. In another implementation, units for the values of the KPIs are displayed. In one implementation, when one of lanes 5302, 5304, 5306 is selected (e.g., by hovering the cursor over the lane) a maximum and a minimum values reflected by the graphical visualization for a corresponding KPI during the current time range are displayed adjacent to visual indicator 5514. For example, in lane 5304, a maximum value of “200” is displayed and a minimum value of “0” is displayed adjacent to visual indicator 5514. This indicates that the highest value of the KPI corresponding to the graphical visualization in lane 5304 during the time period represented by time axis 5410 is “200” and the lowest value during the same time period is “0.” In other implementations, the maximum and minimum values may be displayed for all lanes, regardless of whether they are selected, or may not be displayed for any lanes.

In one implementation, visual interface 5300 may include an indication when the values for a KPI reach one of the



63

predefined KPI thresholds. As discussed above, during the creation of a KPI, the user may define one or more states for the KPI. The states may have corresponding visual characteristics such as colors (e.g., red, yellow, green). In one implementation, the graph color of the graphical visualization may correspond to the color defined for the various states. For example, if the graphical visualization is a line graph, the line may have different colors for values representing different states of the KPI. In another implementation, the current value of a selected lane displayed by visual indicator **5514** may change color to correspond to the colors defined for the various states of the KPI. In another implementation, the values of all lanes displayed by visual indicator **5514** may change color based on the state, regardless of which lane is currently selected. In another implementation, there may be a line or bar running parallel to at least one of lanes **5302**, **5304**, **5306** that is colored according to the colors defined for the various KPI states when the value of the corresponding KPI reaches or passes a defined threshold causing the KPI to change states. In yet another implementation, there may be horizontal lines running along the length of at least one lane to indicate where the thresholds defining different KPI states are located on the vertical axis of the lane. In other implementations, the thresholds may be indicated in visual interface **5300** in some other manner.

FIG. **55B** is a flow diagram of an implementation of a method for inspecting graphical visualizations of KPI values along a time-based graph lane, in accordance with one or more implementations of the present disclosure. At block **5501**, the computing machine determines a point in time corresponding to the current position of lane inspector **5514**. The lane inspector **5514** may be user manipulable such that it may be moved along time axis **5410** to different points in time. For each KPI dataset represented by a graphical visualization in the visual interface, at block **5503**, the computing machine determines a KPI value corresponding to the determined point in time. In addition, at block **5505**, the computing machine determines a state of the KPI at the determined point in time, based on the determined value and the defined KPI thresholds. The determined state may include, for example, a critical state, a warning state, a normal state, etc. At block **5507**, the computing device determines the visual characteristics of the determined state, such as a color (e.g., red, yellow, green) associated with the determined state.

At block **5509**, the computing machine displays the determined value adjacent to lane inspector **5514** for each of the graphical visualizations in the visual interface. In the example illustrated in FIG. **55A**, the value of the graphical visualization in lane **5302** is "3.65," the value of the graphical visualization in lane **5304** is "60," and the value of the graphical visualization in lane **5306** is "0." If the lane inspector **5514** is moved to a new position representing a different time, the operations at blocks **5501-5509** may be repeated.

At block **5511**, the computing machine receives a selection of one of the lanes or graphical visualizations within a lane in the visual interface. In one implementation, one of graph lanes **5302**, **5304**, **5306** is selected by hovering the cursor over the lane. At block **5513**, the computing machine determines the maximum and minimum values of the KPI dataset associated with the selected lane. In one implementation, a search query associated with the KPI is executed for a selected period of time. The results of the query return a dataset of KPI values, as shown in FIG. **45**. The maximum and minimum values from this dataset can be determined. At block **5515**, the computing machine displays the maximum and minimum values adjacent to lane inspector **5515**. For example, in lane

64

**5304**, a maximum value of "200" is displayed and a minimum value of "0" is displayed adjacent to lane inspector **5514**.

FIG. **56** illustrates an example of a visual interface **5300** displaying graphical visualizations of KPI values along time-based graph lanes with options for editing the graphical visualizations, in accordance with one or more implementations of the present disclosure. In one implementation, when one of graph lanes **5302**, **5304**, **5306** is selected (e.g., by hovering the cursor over the lane), a GUI element such as a gear icon **5616** is displayed in association with the selected lane **5306**. When user interaction with gear icon **5616** is detected, a drop down menu **5618** may be displayed. Drop down menu **5618** may include one or more user selectable options including, for example, "Edit Lane," "Delete Lane," "Open in Search," or other options. Selection of one of these options may cause display of a graphical interface to allow the user to edit the graphical visualization in the associated lane **5306**, delete the lane **5306** from the visual interface **5300**, or display the underlying data (e.g., events, machine data) from which the KPI values of the associated graphical visualization are derived. Additional details with respect to editing the graphical visualization are described below in connection with FIG. **57**. When another lane is selected, a corresponding gear icon, or other indicator, may be displayed for the selected lane. In one implementation, a gear icon **5616** is only displayed when the corresponding lane **5306** is selected, and hidden from view when the lane is not selected.

FIG. **57** illustrates an example of a GUI **5700** for editing a graphical visualization of KPI values along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, in response to the selection of the "Edit Lane" option in drop down menu **5618**, the system presents GUI **5700** in order to edit the corresponding graphical visualization.

In one implementation, GUI **5700** can receive user input for a number of input fields **5702**, **5704**, **5712**, selections from drop down menus **5706**, **5708**, or selection of selection buttons **5710** or link **5714**. In one implementation, input field **5702** can be used to edit the title for the graphical visualization. Input field **5204** may be used to edit the subtitle or description of the graphical visualization. In one implementation drop down menu **5706** can be used to edit the graph style, and drop down menu **5708** can be used to edit the graph color for the graphical visualization. For example, upon selection of drop down menu **5708**, a number of available colors may be displayed for selection by the user. Upon selection of a color, the corresponding graphical visualization may be displayed in the selected color. In one implementation, no two graphical visualizations in the same visual interface may have the same color. Accordingly, the available colors displayed for selection may not include any colors already used for other graphical visualizations. In one implementation, the color of a graphical visualization may be determined automatically according to the colors associated with defined thresholds for the corresponding KPI. In such an implementation, the user may not be allowed to edit the graph color in drop down menu **5708**.

Selection buttons **5710** may be used to edit a search source for the graphical visualization. In the illustrated implementation, an "Ad Hoc" search source has been selected. In response, an input field **5712** may display a user-input search query. The search query may include search criteria (e.g., keywords, field/value pairs) that produce a dataset or a search result of events or other data that satisfy the search criteria. In one implementation, a user may edit the search query by making changes, additions, or deletions, to the search query

65

displayed in input field **5712**. The Ad Hoc search query may be executed to generate a dataset of values that can be plotted over the time range as a graphical visualization (e.g., as shown in visual interface **5300**). Selection of link **5714** may indicate that the user wants to execute the search query in input field **5712**. Upon the editing of data and/or the selection menu items, the selection of button **5716** may indicate that the editing of the graphical visualization is complete.

FIG. **58** illustrates an example of a GUI **5700** for editing a graph style of a graphical visualization of KPI values along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, drop down menu **5706** can be used to edit the graph style of the graphical visualization. For example, upon selection of drop down menu **5706**, a list **5806** of available graph types may be displayed for selection by the user. In one implementation, the available graph types include a line graph, an area graph, or a column graph. In other implementations, additional graph types may include a bar chart, a plot graph, a bubble chart, a heat map, or other graph types. Upon selection of a graph type, the corresponding graphical visualization may be displayed in the selected graph type. In one implementation, each graphical visualization on the visual interface has the same graph type. Accordingly, when the graph type of one graphical visualization is changed, the graph type of each remaining graphical visualization in the visual interface is automatically changed to the same graph type. In another implementation, each graphical visualization in the visual interface may have a different graph type. In one implementation, the graph type of a graphical visualization may be determined automatically based on the corresponding KPI or service. In such an implementation, the user may not be allowed to edit the graph type in drop down menu **5706**.

FIG. **59** illustrates an example of a GUI **5700** for selecting the KPI corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, selection buttons **5710** may be used to edit a search source for the graphical visualization. In the illustrated implementation, the “KPI” search source has been selected. In response, drop down menus **5912**, **5914** and input field **5916** may be displayed. Drop down menu **5912** may be used to select a service, the performance of which will be represented by the graphical visualization. Upon selection, drop down menu **5912** may display a list of available services. Drop down menu **5914** may be used to select the KPI that indicates an aspect of how the selected service is performing. Upon selection, drop down menu **5914** may display a list of available KPIs. Input field **5916** may display a search query corresponding to the selected KPI. The search query may derive one or more values from machine data pertaining to one or more entities providing a service. In one implementation, a user may edit the search query by making changes, additions, or deletions, to the search displayed in input field **5916**. Selection of link **5918** may indicate that the user wants to execute the search query in input field **5916**.

FIG. **60** illustrates an example of a GUI **5700** for selecting a data model corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, selection buttons **5710** may be used to edit a search source for the graphical visualization. In the illustrated implementation, the “Data Model” search source has been selected. In response, drop down menus **6012**, **6014** and input fields **6016**, **6018** may be displayed. Drop down menu **6012** may be used to select a data model on which the

66

graphical visualization will be based. Upon selection, drop down menu **6012** may display a list of available data models. Additional details with respect to selection of a data model are described below in connection with FIG. **61**. Drop down menu **6014** may be used to select a statistical function for the data model. Upon selection, drop down menu **6014** may display a list of available functions. Additional details with respect to selection of a data model function are described below in connection with FIG. **62**. Input field **6016** may display a “Where clause” that can be used to further refine the search associated with the selected data model and displayed in input field **6018**. The where clause may include, for example the WHERE command followed by a key/value pair (e.g., WHERE host=Vulcan). In one implementation, “host” is a field name and “Vulcan” is a value stored in the field “host.” The WHERE command may further filter the results of the search query associated with the selected data model to only return data that is associated with the host name “Vulcan.” As a result, the search can filter results based on a particular entity or entities that provide a service. In one implementation, a user may also edit the search query by making changes, additions, or deletions, to the search displayed in input field **6018**. The data model search query may be executed to generate a dataset of values that can be plotted over the time range as a graphical visualization (e.g., as shown in visual interface **5300**). Selection of link **6020** may indicate that the user wants to execute the search query in input field **6018**.

FIG. **61** illustrates an example of a GUI **6100** for selecting a data model corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, upon selection of drop down menu **6012**, GUI **6100** is displayed. GUI **6100** allows for the selection and configuration of a data model to be used as the search source for the graphical visualization. In GUI **6100**, a user may select an existing data model from drop down menu **6102**. Additionally, a user may select one of objects **6104** of the data model. In one implementation, an object is a search that defines one or more events. The data model may be a grouping of objects that are related. Furthermore, a user may select one of the fields **6106** to derive one or more values for the graph. Additional details regarding data models are provided below.

FIG. **62** illustrates an example of a GUI **5700** for editing a statistical function for a data model corresponding to a graphical visualization along a time-based graph lane in a visual interface, in accordance with one or more implementations of the present disclosure. In one implementation, drop down menu **6014** may be used to select statistical function for the data model. For example, upon selection of drop down menu **6014**, a list **6214** of available statistical functions may be displayed for selection by the user. In one implementation, the available statistical functions include average, count, distinct count, maximum, minimum, sum, standard deviation, median or other operations. The selected statistical function may be used to produce one or more values for display as the graphical visualization. In one implementation, the available statistical functions may be dependent on the data type of the selected field from fields **6106** in GUI **6100**. For example, when the selected field has a numerical data type, any of the above listed statistical functions may be available. When the selected field has a string data type, however, the only available operations may be count and distinct count, as the arithmetic operations cannot be performed on a string data type. In one implementation, the statistical function may be determined automatically based on the corresponding data model.

67

In such an implementation, the user may not be allowed to edit the statistical function in drop down menu **5214**.

FIG. **63** illustrates an example of a GUI **6300** for selecting a time range that graphical visualizations along a time-based graph lane in a visual interface should cover, in accordance with one or more implementations of the present disclosure. In one implementation, drop down menu **5308** may be used to select a time range for the graphical visualizations in the visual interface **5300** of FIG. **53**. For example, upon selection of drop down menu **5308**, a GUI **6300** for selection of the time range may be displayed. In one implementation, the time range selection options may include a real-time period **6302**, a relative time period **6304** or some other time period **6306**. For real-time execution, the time range for machine data can be a real-time period **6302** (e.g., 30-second window, 1-minute window, 1-hour window, etc.) from the execution time (e.g., each time the query is executed, the events with timestamps within the specified time window from the query execution time will be used). In real-time execution, a search query associated with the KPI may be continually executed (or periodically executed at a relatively short period (e.g., 1 second)) to continually show a graphical visualization reflecting KPI values from the last one hour (or other real-time period) of time. Thus, if the 1 hour window initially covers from 12 pm to 1 pm, at 1:30, the 1 hour window may cover from 12:30 pm to 1:30 pm. In other words, the time period may be considered a rolling time period, as it constantly changes as time moves forward. For relative execution, the relative time period **6304** can be historical (e.g., yesterday, previous week, etc.) or based on a specified time window from the request time or scheduled time (e.g., last 15 minutes, last 4 hours, etc.). For example, the historical time range “Yesterday” can be selected for relative execution. In another example, the window time range “Last 15 minutes” can be selected for relative execution. In relative execution, the search query associated with the KPI may only be executed upon a request for updated values from the user. Thus, if the 1 hour window covers from 12 pm to 1 pm, that time period will not change until the user requests an update, at which point the most recent 1 hour of values will be displayed. In one implementation, the other time period may include, for example, all of the time where KPI values are available for the corresponding service. Additional time range options may allow the user to specify a particular date or time range over which the KPI values are to be displayed as graphical visualizations.

FIG. **64A** illustrates an example of a visual interface **5300** for selecting a subset of a time range that graphical visualizations along a time-based graph lane in a visual interface cover, in accordance with one or more implementations of the present disclosure. In one implementation, visual indicator **5514** may be used to select a subset **6402** of the time range represented by time axis **5410**, and the corresponding portions of the graphical visualizations in lanes **5302**, **5304**, **5306**. In one implementation, a user may use a mouse or other pointing device to position visual indicator **5514** at a starting position along time axis **5410**, then click and drag to select the desired subset **6402**. In one embodiment, the selected subset **6402** is shown as shaded in the visual interface **5300**. In another implementation, all areas except the selected subset **6402** are shown as shaded. The selection of subset **6402** may be an indication that the user wishes to more closely inspect the KPI values of the graphical visualizations during the time period represented by the subset **6402**. As a result, in response to the selection, the subset **6402** may be emphasized, enlarged, or zoomed in upon to allow closer inspection.

FIG. **64B** is a flow diagram of an implementation of a method for enhancing a view of a subset a subset of a time

68

range for a time-based graph lane, in accordance with one or more implementations of the present disclosure. At block **6401**, the computing device determines a new time range based on the positions of lane inspector **5514**. In one implementation, lane inspector **5514** may be used to select a subset **6402** of the time range represented by time axis **5410**, and the corresponding portions of the graphical visualizations in lanes **5302**, **5304**, **5306**. At block **6403**, the computing device identifies a subset of values of each KPI that correspond to the new time range. In one embodiment, each value in the KPI dataset may have a corresponding time value or timestamp. Thus, the computing device can filter the dataset to identify values with a timestamp included in the selected subset of the time range.

At block **6405**, the computing device determines the maximum and minimum values in the selected subset of values for each KPI, and at block **6407** adjusts the time axis of the lanes in the graphical visualization to reflect the new time range. In one implementation, the subset **6402** is expanded to fill the entire length or nearly the entire length of graph lanes **5302**, **5304**, **5306**. The horizontal axis of each lane may be scaled according to the selected subset **6402**. At block **6409**, the computing device adjusts the height of the lanes based on the new maximum and minimum values. In one implementation, the vertical axis of each lane is scaled according to the maximum and minimum values reflected by the graphical visualization for a corresponding KPI during the selected subset **6402**. At block **6411**, the computing device modifies the graphs based on the subsets of values and calibrates the graphs to the same time scale based on the new time range. Additional details are described with respect to FIG. **65**.

FIG. **65** illustrates an example of a visual interface displaying graphical visualizations of KPI values along time-based graph lanes for a selected subset of a time range, in accordance with one or more implementations of the present disclosure. In response to the selection of subset **6402** using visual indicator **5514**, the system may recalculate the time range that the graphical visualizations in graph lanes **5302**, **5304**, **5306** should cover. In one implementation, the subset **6402** is expanded to fill the entire length or nearly the entire length of graph lanes **5302**, **5304**, **5306**. The horizontal axis of each lane is scaled according to the selected subset **6402** and the vertical axis of each lane is scaled according to the maximum and minimum values reflected by the graphical visualization for a corresponding KPI during the selected subset **6402**. In one implementation, the maximum value appears at or near the top of the lane and the minimum value appears at or near the bottom of the lane. The intermediate values between the maximum and minimum may be scaled accordingly.

In one implementation, time access **5410** is updated according to the selected subset **6402**. The time axis **5410** may include an indication of the amount of time represented by the time scale (e.g., “Viewport: 5 m” indicating that the graphical visualizations in graph lanes **5302**, **5304**, **5306** display KPI values for a time range of five minutes), and an indication of the actual time of day represented by the original time scale (e.g., “12:30, 12:45, 01 PM, 01:15”). In one implementation, a bar running parallel to the time lanes including the indication of the amount of time represented by the time scale (e.g., “Viewport: 1 h 1 m”) is highlighted for a proportional subset of the length of time axis **5410** and only in a location along time axis **5410** corresponding to the subset. In the illustrated embodiment, the highlighted portion of the horizontal bar indicates that the selected subset **6402** occurs sometime between “01 PM” and “01:15.” In one implementation, at least a portion of the time axis **5410** is displayed

above the graph lanes **5302**, **5304**, **5306** as well. This portion of the time axis indicates the actual time of day represented by the selected subset **6402** (e.g., “01:05, 01:06, 01:07, 01:08, 01:09”). In one implementation, a user may return to the un-zoomed view of the original time period by clicking the non-highlighted portion of the horizontal bar in the time axis **5410**.

FIG. **66** illustrates an example of a visual interface **5300** displaying twin graphical visualizations of KPI values along time-based graph lanes for different periods of time, in accordance with one or more implementations of the present disclosure. In one implementation, each of graph lanes **5302**, **5304**, **5306** has a corresponding twin lane **6602**, **6604**, **6606**. The twin lanes **6602**, **6604**, **6606** may display a second graphical visualization in parallel with the first graphical visualization in graph lanes **5302**, **5304**, **5306**. The KPI values reflected in the second graphical visualization may correspond to the same KPI (or other search source) for a different period of time than the values reflected in the first graphical visualization. In one implementation, a user may add the twin lanes **6602**, **6604**, **6606** by selecting drop down menu **6608**. In one implementation, drop down menu **6608** can be used to select the period of time for the values reflected in the second graphical visualizations. For example, upon selection of drop down menu **6608**, a list **6610** of available time periods may be displayed for selection by the user. In one implementation, the available time periods may include periods of time in the past when KPI data is available for one or more of the graphical visualizations. In one implementation, a twin lane may be created for each of the lanes in the visual interface, and a search query of each KPI can be executed using the specified time range to produce one or more time values for the second graphical visualization of a corresponding KPI. Because the new time range is associated with a different point(s) in time, the machine data or events used by the search query for the second graphical visualization will be different than the machine data that was used by the search query for the original graphical visualization, and therefore the values produced for the second graphical visualization are likely to be different from the values that were produced for the original graphical visualization. In another implementation, a twin lane may be created only for one or more selected lanes in the visual interface, and only search queries of those KPIs can be executed. In one implementation, if past KPI data is not available for the selected time range, no second graphical visualization may be displayed in the twin lane **6606**.

FIG. **67** illustrates an example of a visual interface with a user manipulable visual indicator **5514** spanning across twin graphical visualizations of KPI values along time-based graph lanes for different periods of time, in accordance with one or more implementations of the present disclosure. Visual indicator **5514**, also referred to herein as a “lane inspector,” may include, for example, a line or other indicator that spans across the graph lanes **5302**, **6602**, **5304**, **6604**, **5306**, **6606** at a given point in time along time axis **5410**. The visual indicator **5514** may be user manipulable such that it may be moved along time axis **5410** to different points. For example, visual indicator **5514** may slide back and forth along the lengths of graph lanes and time axis **5410** in response to user input received with a mouse, touchpad, touchscreen, etc.

In one implementation, visual indicator **5514** includes a display of the point in time at which it is currently located both in original lanes **5302**, **5304**, **5306** and twin lanes **6602**, **6604**, **6606**. In the illustrated example, the times associated with visual indicator **5514** are “Thu Sep 4 01:35:34 PM” for the original lanes and “Wed Sep 3 01:35:34 PM” for the twin lanes. Thus, the twin lanes show values of the same KPI from

the same time range on the previous day. In one implementation, visual indicator **5514** further includes a display of a value reflected in each of the graphical visualizations for the different KPIs at the point in time corresponding to the position of visual indicator **5514**. In the illustrated example, the value of the graphical visualization in lane **5302** is “0,” the value of the graphical visualization in lane **6302** is “1.52,” the value of the graphical visualization in lane **5304** is “36,” the value of the graphical visualization in lane **6304** is “31,” the value of the graphical visualization in lane **5306** is “0,” and lane **6306** has no data available. In one implementation, the graphical visualizations in twin lanes **6302**, **6304**, **6306** have the same graph type and a similar graph color as the graphical visualizations in the corresponding graph lanes **5302**, **5304**, **5306**. In another implementation, the second graphical visualizations are configurable such that the user can adjust the graph type and the graph color. In one implementation, rather than being displayed in twin parallel lanes, the second graphical visualizations may be overlaid on top of the original graphical visualizations.

FIG. **68** illustrates an example of a visual interface **5300** displaying a graph lane **6806** with inventory information for a service or entities reflected by KPI values, in accordance with one or more implementations of the present disclosure. In one implementation, an additional lane **6806** is displayed in parallel to at least one of graph lanes **6802** and **6804**. Graph lanes **6802** and **6804** may be similar to graph lanes **5302**, **5304**, **5306** described above, such that they may display graphical visualizations of corresponding KPI values. Additional lane **6806**, however, may be a different type of lane, which does not display graphical visualizations. In one implementation, additional lane **6806** may display inventory information for the service or for the one or more entities providing the service reflected by the KPI corresponding to the graphical visualization in the adjacent lane **6804**. The additional lane **6806** may include textual information, or other non-graphical information. The inventory information may include information about the service or the entities providing the service, such as an identifier of the entities (e.g., a host name, server name), a location of the entities (e.g., rack number, data center name), etc. In one implementation, the inventory information displayed in lane **6806** may be populated from information provided during the entity definition process. In one embodiment, the inventory information displayed in additional lane **6806** may change according to the position of visual indicator **5514** along time axis **5410**. When the inventory information is time stamped, or otherwise is associated with a time value, the inventory information may be different at different points in time. Accordingly, in one implementation, the inventory information available at the time associated with the position of visual indicator **5514** may be displayed in additional lane **6806**. In one implementation, additional lane **6806** may be continually associated with an adjacent lane **6804**, such that if the lanes in visual interface **5300** are reordered, additional lane **6806** remains adjacent to lane **6804** despite the reordering.

FIG. **69** illustrates an example of a visual interface **5300** displaying a graph lane with notable events occurring during a timer period covered by graphical visualization of KPI values, in accordance with one or more implementations of the present disclosure. In one implementation, an additional lane **6908** is displayed in parallel to at least one of graph lanes **6902**, **6904**, **6906**. Graph lanes **6902**, **6904**, **6906** may be similar to graph lanes **5302**, **5304**, **5306** described above, such that they may display graphical visualizations of corresponding KPI values. Additional lane **6908**, however, may be a different type of lane designed to display indications of the

occurrences of notable events. “Notable events” are system occurrences that may be likely to indicate a security threat or operational problem. These notable events can be detected in a number of ways: (1) an analyst can notice a correlation in the data and can manually identify a corresponding group of one or more events as “notable;” or (2) an analyst can define a “correlation search” specifying criteria for a notable event, and every time one or more events satisfy the criteria, the application can indicate that the one or more events are notable. An analyst can alternatively select a pre-defined correlation search provided by the application. Note that correlation searches can be run continuously or at regular intervals (e.g., every hour) to search for notable events. Upon detection, notable events can be stored in a dedicated “notable events index,” which can be subsequently accessed to generate various visualizations containing security-related information.

In one implementation, the notable events occurring during the period of time represented by time axis **5410** are displayed as flags **6910** or bubbles in a bubble chart in additional lane **6908**. The flags **6910** may be located at a position along time axis **5410** corresponding to when the notable event occurred. In one implementation, the flags **6910** may be color coded to vindicate the severity or importance of the notable event. In one implementation, when one of the flags **6910** is selected (e.g., by clicking on the flag or hovering the cursor over the flag), a description of the notable event may be displayed. As illustrated in FIG. 69, the description **6912** may be displayed in a horizontal bar along the bottom of lane **6908**. In another implementation, as illustrated in FIG. 70, the description **7012** may be displayed adjacent to the selected flag **6910**. In one implementation, user-manipulable visual indicator **5514** may be used to select a particular flag **6910**. For example, when visual indicator **5514** is slid along the length of lane **6908**, a description **7012** of a corresponding notable event at the same time may be displayed.

In some implementations, search queries for KPIs and correlation searches can derive values using a late binding schema that the search queries apply to machine data. Late binding schema is described in greater detail below. The systems and methods described herein above may be employed by various data processing systems, e.g., data aggregation and analysis systems. In various illustrative examples, the data processing system may be represented by the SPLUNK® ENTERPRISE system produced by Splunk Inc. of San Francisco, Calif., to store and process performance data.

### 1.1 Overview

Modern data centers often comprise thousands of host computer systems that operate collectively to service requests from even larger numbers of remote clients. During operation, these data centers generate significant volumes of performance data and diagnostic information that can be analyzed to quickly diagnose performance problems. In order to reduce the size of this performance data, the data is typically pre-processed prior to being stored based on anticipated data-analysis needs. For example, pre-specified data items can be extracted from the performance data and stored in a database to facilitate efficient retrieval and analysis at search time. However, the rest of the performance data is not saved and is essentially discarded during pre-processing. As storage capacity becomes progressively cheaper and more plentiful, there are fewer incentives to discard this performance data and many reasons to keep it.

This plentiful storage capacity is presently making it feasible to store massive quantities of minimally processed performance data at “ingestion time” for later retrieval and

analysis at “search time.” Note that performing the analysis operations at search time provides greater flexibility because it enables an analyst to search all of the performance data, instead of searching pre-specified data items that were stored at ingestion time. This enables the analyst to investigate different implementations of the performance data instead of being confined to the pre-specified set of data items that were selected at ingestion time.

However, analyzing massive quantities of heterogeneous performance data at search time can be a challenging task. A data center may generate heterogeneous performance data from thousands of different components, which can collectively generate tremendous volumes of performance data that can be time-consuming to analyze. For example, this performance data can include data from system logs, network packet data, sensor data, and data generated by various applications. Also, the unstructured nature of much of this performance data can pose additional challenges because of the difficulty of applying semantic meaning to unstructured data, and the difficulty of indexing and querying unstructured data using traditional database systems.

These challenges can be addressed by using an event-based system, such as the SPLUNK® ENTERPRISE system produced by Splunk Inc. of San Francisco, Calif., to store and process performance data. The SPLUNK® ENTERPRISE system is the leading platform for providing real-time operational intelligence that enables organizations to collect, index, and harness machine-generated data from various websites, applications, servers, networks, and mobile devices that power their businesses. The SPLUNK® ENTERPRISE system is particularly useful for analyzing unstructured performance data, which is commonly found in system log files. Although many of the techniques described herein are explained with reference to the SPLUNK® ENTERPRISE system, the techniques are also applicable to other types of data server systems.

In the SPLUNK® ENTERPRISE system, performance data is stored as “events,” wherein each event comprises a collection of performance data and/or diagnostic information that is generated by a computer system and is correlated with a specific point in time. Events can be derived from “time series data,” wherein time series data comprises a sequence of data points (e.g., performance measurements from a computer system) that are associated with successive points in time and are typically spaced at uniform time intervals. Events can also be derived from “structured” or “unstructured” data. Structured data has a predefined format, wherein specific data items with specific data formats reside at predefined locations in the data. For example, structured data can include data items stored in fields in a database table. In contrast, unstructured data does not have a predefined format. This means that unstructured data can comprise various data items having different data types that can reside at different locations. For example, when the data source is an operating system log, an event can include one or more lines from the operating system log containing raw data that includes different types of performance and diagnostic information associated with a specific point in time. Examples of data sources from which an event may be derived include, but are not limited to: web servers; application servers; databases; firewalls; routers; operating systems; and software applications that execute on computer systems, mobile devices, and sensors. The data generated by such data sources can be produced in various forms including, for example and without limitation, server log files, activity log files, configuration files, messages, network packet data, performance measurements and sensor measurements. An event typically includes a

timestamp that may be derived from the raw data in the event, or may be determined through interpolation between temporally proximate events having known timestamps.

The SPLUNK® ENTERPRISE system also facilitates using a flexible schema to specify how to extract information from the event data, wherein the flexible schema may be developed and redefined as needed. Note that a flexible schema may be applied to event data “on the fly,” when it is needed (e.g., at search time), rather than at ingestion time of the data as in traditional database systems. Because the schema is not applied to event data until it is needed (e.g., at search time), it is referred to as a “late-binding schema.”

During operation, the SPLUNK® ENTERPRISE system starts with raw data, which can include unstructured data, machine data, performance measurements or other time-series data, such as data obtained from weblogs, syslogs, or sensor readings. It divides this raw data into “portions,” and optionally transforms the data to produce timestamped events. The system stores the timestamped events in a data store, and enables a user to run queries against the data store to retrieve events that meet specified criteria, such as containing certain keywords or having specific values in defined fields. Note that the term “field” refers to a location in the event data containing a value for a specific data item.

As noted above, the SPLUNK® ENTERPRISE system facilitates using a late-binding schema while performing queries on events. A late-binding schema specifies “extraction rules” that are applied to data in the events to extract values for specific fields. More specifically, the extraction rules for a field can include one or more instructions that specify how to extract a value for the field from the event data. An extraction rule can generally include any type of instruction for extracting values from data in events. In some cases, an extraction rule comprises a regular expression, in which case the rule is referred to as a “regex rule.”

In contrast to a conventional schema for a database system, a late-binding schema is not defined at data ingestion time. Instead, the late-binding schema can be developed on an ongoing basis until the time a query is actually executed. This means that extraction rules for the fields in a query may be provided in the query itself, or may be located during execution of the query. Hence, as an analyst learns more about the data in the events, the analyst can continue to refine the late-binding schema by adding new fields, deleting fields, or changing the field extraction rules until the next time the schema is used by a query. Because the SPLUNK® ENTERPRISE system maintains the underlying raw data and provides a late-binding schema for searching the raw data, it enables an analyst to investigate questions that arise as the analyst learns more about the events.

In the SPLUNK® ENTERPRISE system, a field extractor may be configured to automatically generate extraction rules for certain fields in the events when the events are being created, indexed, or stored, or possibly at a later time. Alternatively, a user may manually define extraction rules for fields using a variety of techniques.

Also, a number of “default fields” that specify metadata about the events rather than data in the events themselves can be created automatically. For example, such default fields can specify: a timestamp for the event data; a host from which the event data originated; a source of the event data; and a source type for the event data. These default fields may be determined automatically when the events are created, indexed or stored.

In some embodiments, a common field name may be used to reference two or more fields containing equivalent data items, even though the fields may be associated with different

types of events that possibly have different data formats and different extraction rules. By enabling a common field name to be used to identify equivalent fields from different types of events generated by different data sources, the system facilitates use of a “common information model” (CIM) across the different data sources.

#### 1.2 Data Server System

FIG. 71 presents a block diagram of an exemplary event-processing system 7100, similar to the SPLUNK® ENTERPRISE system. System 7100 includes one or more forwarders 7101 that collect data obtained from a variety of different data sources 7105, and one or more indexers 7102 that store, process, and/or perform operations on this data, wherein each indexer operates on data contained in a specific data store 7103. These forwarders and indexers can comprise separate computer systems in a data center, or may alternatively comprise separate processes executing on various computer systems in a data center.

During operation, the forwarders 7101 identify which indexers 7102 will receive the collected data and then forward the data to the identified indexers. Forwarders 7101 can also perform operations to strip out extraneous data and detect timestamps in the data. The forwarders next determine which indexers 7102 will receive each data item and then forward the data items to the determined indexers 7102.

Note that distributing data across different indexers facilitates parallel processing. This parallel processing can take place at data ingestion time, because multiple indexers can process the incoming data in parallel. The parallel processing can also take place at search time, because multiple indexers can search through the data in parallel.

System 7100 and the processes described below with respect to FIGS. 71-5 are further described in “Exploring Splunk Search Processing Language (SPL) Primer and Cookbook” by David Carasso, CITO Research, 2012, and in “Optimizing Data Analysis With a Semi-Structured Time Series Database” by Ledion Bitincka, Archana Ganapathi, Stephen Sorkin, and Steve Zhang, SLAML, 2010, each of which is hereby incorporated herein by reference in its entirety for all purposes.

#### 1.3 Data Ingestion

FIG. 72 presents a flowchart illustrating how an indexer processes, indexes, and stores data received from forwarders in accordance with the disclosed embodiments. At block 7201, the indexer receives the data from the forwarder. Next, at block 7202, the indexer apportions the data into events. Note that the data can include lines of text that are separated by carriage returns or line breaks and an event may include one or more of these lines. During the apportioning process, the indexer can use heuristic rules to automatically determine the boundaries of the events, which for example coincide with line boundaries. These heuristic rules may be determined based on the source of the data, wherein the indexer can be explicitly informed about the source of the data or can infer the source of the data by examining the data. These heuristic rules can include regular expression-based rules or delimiter-based rules for determining event boundaries, wherein the event boundaries may be indicated by predefined characters or character strings. These predefined characters may include punctuation marks or other special characters including, for example, carriage returns, tabs, spaces or line breaks. In some cases, a user can fine-tune or configure the rules that the indexers use to determine event boundaries in order to adapt the rules to the user’s specific requirements.

Next, the indexer determines a timestamp for each event at block 7203. As mentioned above, these timestamps can be determined by extracting the time directly from data in the

75

event, or by interpolating the time based on timestamps from temporally proximate events. In some cases, a timestamp can be determined based on the time the data was received or generated. The indexer subsequently associates the determined timestamp with each event at block 7204, for example by storing the timestamp as metadata for each event.

Then, the system can apply transformations to data to be included in events at block 7205. For log data, such transformations can include removing a portion of an event (e.g., a portion used to define event boundaries, extraneous text, characters, etc.) or removing redundant portions of an event. Note that a user can specify portions to be removed using a regular expression or any other possible technique.

Next, a keyword index can optionally be generated to facilitate fast keyword searching for events. To build a keyword index, the indexer first identifies a set of keywords in block 7206. Then, at block 7207 the indexer includes the identified keywords in an index, which associates each stored keyword with references to events containing that keyword (or to locations within events where that keyword is located). When an indexer subsequently receives a keyword-based query, the indexer can access the keyword index to quickly identify events containing the keyword.

In some embodiments, the keyword index may include entries for name-value pairs found in events, wherein a name-value pair can include a pair of keywords connected by a symbol, such as an equals sign or colon. In this way, events containing these name-value pairs can be quickly located. In some embodiments, fields can automatically be generated for some or all of the name-value pairs at the time of indexing. For example, if the string "dest=10.0.1.2" is found in an event, a field named "dest" may be created for the event, and assigned a value of "10.0.1.2."

Finally, the indexer stores the events in a data store at block 7208, wherein a timestamp can be stored with each event to facilitate searching for events based on a time range. In some cases, the stored events are organized into a plurality of buckets, wherein each bucket stores events associated with a specific time range. This not only improves time-based searches, but it also allows events with recent timestamps that may have a higher likelihood of being accessed to be stored in faster memory to facilitate faster retrieval. For example, a bucket containing the most recent events can be stored as flash memory instead of on hard disk.

Each indexer 7102 is responsible for storing and searching a subset of the events contained in a corresponding data store 7103. By distributing events among the indexers and data stores, the indexers can analyze events for a query in parallel, for example using map-reduce techniques, wherein each indexer returns partial responses for a subset of events to a search head that combines the results to produce an answer for the query. By storing events in buckets for specific time ranges, an indexer may further optimize searching by looking only in buckets for time ranges that are relevant to a query.

Moreover, events and buckets can also be replicated across different indexers and data stores to facilitate high availability and disaster recovery as is described in U.S. patent application Ser. No. 14/266,812 filed on 30 Apr. 2014, and in U.S. patent application Ser. No. 14/266,817 also filed on 30 Apr. 2014.

#### 1.4 Query Processing

FIG. 73 presents a flowchart illustrating how a search head and indexers perform a search query in accordance with the disclosed embodiments. At the start of this process, a search head receives a search query from a client at block 7301. Next, at block 7302, the search head analyzes the search query to determine what portions can be delegated to indexers

76

and what portions need to be executed locally by the search head. At block 7303, the search head distributes the determined portions of the query to the indexers. Note that commands that operate on single events can be trivially delegated to the indexers, while commands that involve events from multiple indexers are harder to delegate.

Then, at block 7304, the indexers to which the query was distributed search their data stores for events that are responsive to the query. To determine which events are responsive to the query, the indexer searches for events that match the criteria specified in the query. This criteria can include matching keywords or specific values for certain fields. In a query that uses a late-binding schema, the searching operations in block 7304 may involve using the late-binding scheme to extract values for specified fields from events at the time the query is processed. Next, the indexers can either send the relevant events back to the search head, or use the events to calculate a partial result, and send the partial result back to the search head.

Finally, at block 7305, the search head combines the partial results and/or events received from the indexers to produce a final result for the query. This final result can comprise different types of data depending upon what the query is asking for. For example, the final results can include a listing of matching events returned by the query, or some type of visualization of data from the returned events. In another example, the final result can include one or more calculated values derived from the matching events.

Moreover, the results generated by system 7100 can be returned to a client using different techniques. For example, one technique streams results back to a client in real-time as they are identified. Another technique waits to report results to the client until a complete set of results is ready to return to the client. Yet another technique streams interim results back to the client in real-time until a complete set of results is ready, and then returns the complete set of results to the client. In another technique, certain results are stored as "search jobs," and the client may subsequently retrieve the results by referencing the search jobs.

The search head can also perform various operations to make the search more efficient. For example, before the search head starts executing a query, the search head can determine a time range for the query and a set of common keywords that all matching events must include. Next, the search head can use these parameters to query the indexers to obtain a superset of the eventual results. Then, during a filtering stage, the search head can perform field-extraction operations on the superset to produce a reduced set of search results.

#### 1.5 Field Extraction

FIG. 74A presents a block diagram illustrating how fields can be extracted during query processing in accordance with the disclosed embodiments. At the start of this process, a search query 7402 is received at a query processor 7404. Query processor 7404 includes various mechanisms for processing a query, wherein these mechanisms can reside in a search head 7104 and/or an indexer 7102. Note that the exemplary search query 7402 illustrated in FIG. 74A is expressed in Search Processing Language (SPL), which is used in conjunction with the SPLUNK® ENTERPRISE system. SPL is a pipelined search language in which a set of inputs is operated on by a first command in a command line, and then a subsequent command following the pipe symbol "|" operates on the results produced by the first command, and so on for additional commands. Search query 7402 can also be expressed in other query languages, such as the Structured Query Language ("SQL") or any suitable query language.



Upon receiving search query 7402, query processor 7404 sees that search query 7402 includes two fields “IP” and “target.” Query processor 7404 also determines that the values for the “IP” and “target” fields have not already been extracted from events in data store 7414, and consequently determines that query processor 7404 needs to use extraction rules to extract values for the fields. Hence, query processor 7404 performs a lookup for the extraction rules in a rule base 7406, wherein rule base 7406 maps field names to corresponding extraction rules and obtains extraction rules 7408-7409, wherein extraction rule 7408 specifies how to extract a value for the “IP” field from an event, and extraction rule 7409 specifies how to extract a value for the “target” field from an event. As is illustrated in FIG. 74A, extraction rules 7408-7409 can comprise regular expressions that specify how to extract values for the relevant fields. Such regular-expression-based extraction rules are also referred to as “regex rules.” In addition to specifying how to extract field values, the extraction rules may also include instructions for deriving a field value by performing a function on a character string or value retrieved by the extraction rule. For example, a transformation rule may truncate a character string, or convert the character string into a different data format. In some cases, the query itself can specify one or more extraction rules.

Next, query processor 7404 sends extraction rules 7408-7409 to a field extractor 7412, which applies extraction rules 7408-7409 to events 7416-7418 in a data store 7414. Note that data store 7414 can include one or more data stores, and extraction rules 7408-7409 can be applied to large numbers of events in data store 7414, and are not meant to be limited to the three events 7416-7418 illustrated in FIG. 74A. Moreover, the query processor 7404 can instruct field extractor 7412 to apply the extraction rules to all the events in a data store 7414, or to a subset of the events that have been filtered based on some criteria.

Next, field extractor 7412 applies extraction rule 7408 for the first command “Search IP=“10\*” to events in data store 7414 including events 7416-7418. Extraction rule 7408 is used to extract values for the IP address field from events in data store 7414 by looking for a pattern of one or more digits, followed by a period, followed again by one or more digits, followed by another period, followed again by one or more digits, followed by another period, and followed again by one or more digits. Next, field extractor 7412 returns field values 7420 to query processor 7404, which uses the criterion IP=“10\*” to look for IP addresses that start with “10”. Note that events 7416 and 7417 match this criterion, but event 7418 does not, so the result set for the first command is events 7416-7417.

Query processor 7404 then sends events 7416-717 to the next command “stats count target.” To process this command, query processor 7404 causes field extractor 7412 to apply extraction rule 7409 to events 7416-7417. Extraction rule 7409 is used to extract values for the target field for events 7416-7417 by skipping the first four commas in events 7416-7417, and then extracting all of the following characters until a comma or period is reached. Next, field extractor 7412 returns field values 7421 to query processor 7404, which executes the command “stats count target” to count the number of unique values contained in the target fields, which in this example produces the value “2” that is returned as a final result 7422 for the query.

Note that query results can be returned to a client, a search head, or any other system component for further processing. In general, query results may include: a set of one or more events; a set of one or more values obtained from the events; a subset of the values; statistics calculated based on the val-

ues; a report containing the values; or a visualization, such as a graph or chart, generated from the values.

#### 1.5.1 Data Models

Creating queries requires knowledge of the fields that are included in the events being searched, as well as knowledge of the query processing language used for the queries. While a data analyst may possess domain understanding of underlying data and knowledge of the query processing language, an end user responsible for creating reports at a company (e.g., a marketing specialist) may not have such expertise. In order to assist end users, implementations of the event-processing system described herein provide data models that simplify the creation of reports and other visualizations.

A data model encapsulates semantic knowledge about certain events. A data model can be composed of one or more objects grouped in a hierarchical manner. In general, the objects included in a data model may be related to each other in some way. In particular, a data model can include a root object and, optionally, one or more child objects that can be linked (either directly or indirectly) to the root object. A root object can be defined by search criteria for a query to produce a certain set of events, and a set of fields that can be exposed to operate on those events. A root object can be a parent of one or more child objects, and any of those child objects can optionally be a parent of one or more additional child objects. Each child object can inherit the search criteria of its parent object and have additional search criteria to further filter out events represented by its parent object. Each child object may also include at least some of the fields of its parent object and optionally additional fields specific to the child object.

FIG. 74B illustrates an example data model structure 7428, in accordance with some implementations. As shown, example data model “Buttercup Games” 7430 includes root object “Purchase Requests” 7432, and child objects “Successful Purchases” 7434 and “Unsuccessful Purchases” 7436.

FIG. 74C illustrates an example definition 7440 of root object 7432 of data model 7430, in accordance with some implementations. As shown, definition 7440 of root object 7432 includes search criteria 7442 and a set of fields 7444. Search criteria 7442 require that a search query produce web access requests that qualify as purchase events. Fields 7444 include inherited fields 7446 which are default fields that specify metadata about the events of the root object 7432. In addition, fields 7444 include extracted fields 7448, whose values can be automatically extracted from the events during search using extraction rules of the late binding schema, and calculated fields 7450, whose values can be automatically determined based on values of other fields extracted from the events. For example, the value of the productName field can be determined based on the value in the productID field (e.g., by searching a lookup table for a product name matching the value of the productID field). In another example, the value of the price field can be calculated based on values of other fields (e.g., by multiplying the price per unit by the number of units).

FIG. 74D illustrates example definitions 7458 and 7460 of child objects 7434 and 7436 respectively, in accordance with some implementations. Definition 7458 of child object 7434 includes search criteria 7462 and a set of fields 7464. Search criteria 7462 inherits search criteria 7442 of the parent object 7432 and includes an additional criterion of “status=200,” which indicates that the search query should produce web access requests that qualify as successful purchase events. Fields 7464 consist of the fields inherited from the parent object 7432.



Definition **7460** of child object **7436** includes search criteria **7470** and a set of fields **7474**. Search criteria **7470** inherits search criteria **7442** of the parent object **7432** and includes an additional criterion of “status!=200,” which indicates that the search query should produce web access requests that qualify as unsuccessful purchase events. Fields **7474** consist of the fields inherited from the parent object **7432**. As shown, child objects **7434** and **7436** include all the fields inherited from the parent object **7432**. In other implementations, child objects may only include some of the fields of the parent object and/or may include additional fields that are not exposed by the parent object.

When creating a report, a user can select an object of a data model to focus on the events represented by the selected object. The user can then view the fields of the data object and request the event-processing system to structure the report based on those fields. For example, the user can request the event-processing system to add some fields to the report, to add calculations based on some fields to the report, to group data in the report based on some fields, etc. The user can also input additional constraints (e.g., specific values and/or mathematical expressions) for some of the fields to further filter out events on which the report should be focused.

#### 1.6 Exemplary Search Screen

FIG. **76A** illustrates an exemplary search screen **7600** in accordance with the disclosed embodiments. Search screen **7600** includes a search bar **7602** that accepts user input in the form of a search string. It also includes a time range picker **7612** that enables the user to specify a time range for the search. For “historical searches” the user can select a specific time range, or alternatively a relative time range, such as “today,” “yesterday” or “last week.” For “real-time searches,” the user can select the size of a preceding time window to search for real-time events. Search screen **7600** also initially displays a “data summary” dialog as is illustrated in FIG. **76B** that enables the user to select different sources for the event data, for example by selecting specific hosts and log files.

After the search is executed, the search screen **7600** can display the results through search results tabs **7604**, wherein search results tabs **7604** includes: an “events tab” that displays various information about events returned by the search; a “statistics tab” that displays statistics about the search results; and a “visualization tab” that displays various visualizations of the search results. The events tab illustrated in FIG. **76A** displays a timeline graph **7605** that graphically illustrates the number of events that occurred in one-hour intervals over the selected time range. It also displays an events list **7608** that enables a user to view the raw data in each of the returned events. It additionally displays a fields sidebar **7606** that includes statistics about occurrences of specific fields in the returned events, including “selected fields” that are pre-selected by the user, and “interesting fields” that are automatically selected by the system based on pre-specified criteria.

#### 1.7 Acceleration Techniques

The above-described system provides significant flexibility by enabling a user to analyze massive quantities of minimally processed performance data “on the fly” at search time instead of storing pre-specified portions of the performance data in a database at ingestion time. This flexibility enables a user to see correlations in the performance data and perform subsequent queries to examine interesting implementations of the performance data that may not have been apparent at ingestion time.

However, performing extraction and analysis operations at search time can involve a large amount of data and require a large number of computational operations, which can cause

considerable delays while processing the queries. Fortunately, a number of acceleration techniques have been developed to speed up analysis operations performed at search time. These techniques include: (1) performing search operations in parallel by formulating a search as a map-reduce computation; (2) using a keyword index; (3) using a high performance analytics store; and (4) accelerating the process of generating reports. These techniques are described in more detail below.

##### 1.7.1 Map-Reduce Technique

To facilitate faster query processing, a query can be structured as a map-reduce computation, wherein the “map” operations are delegated to the indexers, while the corresponding “reduce” operations are performed locally at the search head. For example, FIG. **75** illustrates how a search query **7501** received from a client at search head **7104** can split into two phases, including: (1) a “map phase” comprising subtasks **7502** (e.g., data retrieval or simple filtering) that may be performed in parallel and are “mapped” to indexers **7102** for execution, and (2) a “reduce phase” comprising a merging operation **7503** to be executed by the search head when the results are ultimately collected from the indexers.

During operation, upon receiving search query **7501**, search head **7104** modifies search query **7501** by substituting “stats” with “prestats” to produce search query **7502**, and then distributes search query **7502** to one or more distributed indexers, which are also referred to as “search peers.” Note that search queries may generally specify search criteria or operations to be performed on events that meet the search criteria. Search queries may also specify field names, as well as search criteria for the values in the fields or operations to be performed on the values in the fields. Moreover, the search head may distribute the full search query to the search peers as is illustrated in FIG. **73**, or may alternatively distribute a modified version (e.g., a more restricted version) of the search query to the search peers. In this example, the indexers are responsible for producing the results and sending them to the search head. After the indexers return the results to the search head, the search head performs the merging operations **7503** on the results. Note that by executing the computation in this way, the system effectively distributes the computational operations while minimizing data transfers.

##### 1.7.2 Keyword Index

As described above with reference to the flow charts in FIGS. **72** and **73**, event-processing system **7100** can construct and maintain one or more keyword indices to facilitate rapidly identifying events containing specific keywords. This can greatly speed up the processing of queries involving specific keywords. As mentioned above, to build a keyword index, an indexer first identifies a set of keywords. Then, the indexer includes the identified keywords in an index, which associates each stored keyword with references to events containing that keyword, or to locations within events where that keyword is located. When an indexer subsequently receives a keyword-based query, the indexer can access the keyword index to quickly identify events containing the keyword.

##### 1.7.3 High Performance Analytics Store

To speed up certain types of queries, some embodiments of system **7100** make use of a high performance analytics store, which is referred to as a “summarization table,” that contains entries for specific field-value pairs. Each of these entries keeps track of instances of a specific value in a specific field in the event data and includes references to events containing the specific value in the specific field. For example, an exemplary entry in a summarization table can keep track of occurrences of the value “94107” in a “ZIP code” field of a set of

events, wherein the entry includes references to all of the events that contain the value “94107” in the ZIP code field. This enables the system to quickly process queries that seek to determine how many events have a particular value for a particular field, because the system can examine the entry in the summarization table to count instances of the specific value in the field without having to go through the individual events or do extractions at search time. Also, if the system needs to process all events that have a specific field-value combination, the system can use the references in the summarization table entry to directly access the events to extract further information without having to search all of the events to find the specific field-value combination at search time.

In some embodiments, the system maintains a separate summarization table for each of the above-described time-specific buckets that stores events for a specific time range, wherein a bucket-specific summarization table includes entries for specific field-value combinations that occur in events in the specific bucket. Alternatively, the system can maintain a separate summarization table for each indexer, wherein the indexer-specific summarization table only includes entries for the events in a data store that is managed by the specific indexer.

The summarization table can be populated by running a “collection query” that scans a set of events to find instances of a specific field-value combination, or alternatively instances of all field-value combinations for a specific field. A collection query can be initiated by a user, or can be scheduled to occur automatically at specific time intervals. A collection query can also be automatically launched in response to a query that asks for a specific field-value combination.

In some cases, the summarization tables may not cover all of the events that are relevant to a query. In this case, the system can use the summarization tables to obtain partial results for the events that are covered by summarization tables, but may also have to search through other events that are not covered by the summarization tables to produce additional results. These additional results can then be combined with the partial results to produce a final set of results for the query. This summarization table and associated techniques are described in more detail in U.S. Pat. No. 8,682,925, issued on Mar. 25, 2014.

#### 1.7.4 Accelerating Report Generation

In some embodiments, a data server system such as the SPLUNK® ENTERPRISE system can accelerate the process of periodically generating updated reports based on query results. To accelerate this process, a summarization engine automatically examines the query to determine whether generation of updated reports can be accelerated by creating intermediate summaries. (This is possible if results from preceding time periods can be computed separately and combined to generate an updated report. In some cases, it is not possible to combine such incremental results, for example where a value in the report depends on relationships between events from different time periods.) If reports can be accelerated, the summarization engine periodically generates a summary covering data obtained during a latest non-overlapping time period. For example, where the query seeks events meeting a specified criteria, a summary for the time period includes only events within the time period that meet the specified criteria. Similarly, if the query seeks statistics calculated from the events, such as the number of events that match the specified criteria, then the summary for the time period includes the number of events in the period that match the specified criteria.

In parallel with the creation of the summaries, the summarization engine schedules the periodic updating of the report

associated with the query. During each scheduled report update, the query engine determines whether intermediate summaries have been generated covering portions of the time period covered by the report update. If so, then the report is generated based on the information contained in the summaries. Also, if additional event data has been received and has not yet been summarized, and is required to generate the complete report, the query can be run on this additional event data. Then, the results returned by this query on the additional event data, along with the partial results obtained from the intermediate summaries, can be combined to generate the updated report. This process is repeated each time the report is updated. Alternatively, if the system stores events in buckets covering specific time ranges, then the summaries can be generated on a bucket-by-bucket basis. Note that producing intermediate summaries can save the work involved in re-running the query for previous time periods, so only the newer event data needs to be processed while generating an updated report. These report acceleration techniques are described in more detail in U.S. Pat. No. 8,589,403, issued on Nov. 19, 2013, and U.S. Pat. No. 8,412,696, issued on Apr. 2, 2011.

#### 1.8 Security Features

The SPLUNK® ENTERPRISE platform provides various schemas, dashboards and visualizations that make it easy for developers to create applications to provide additional capabilities. One such application is the SPLUNK® APP FOR ENTERPRISE SECURITY, which performs monitoring and alerting operations and includes analytics to facilitate identifying both known and unknown security threats based on large volumes of data stored by the SPLUNK® ENTERPRISE system. This differs significantly from conventional Security Information and Event Management (SIEM) systems that lack the infrastructure to effectively store and analyze large volumes of security-related event data. Traditional SIEM systems typically use fixed schemas to extract data from pre-defined security-related fields at data ingestion time, wherein the extracted data is typically stored in a relational database. This data extraction process (and associated reduction in data size) that occurs at data ingestion time inevitably hampers future incident investigations, when all of the original data may be needed to determine the root cause of a security issue, or to detect the tiny fingerprints of an impending security threat.

In contrast, the SPLUNK® APP FOR ENTERPRISE SECURITY system stores large volumes of minimally processed security-related data at ingestion time for later retrieval and analysis at search time when a live security threat is being investigated. To facilitate this data retrieval process, the SPLUNK® APP FOR ENTERPRISE SECURITY provides pre-specified schemas for extracting relevant values from the different types of security-related event data, and also enables a user to define such schemas.

The SPLUNK® APP FOR ENTERPRISE SECURITY can process many types of security-related information. In general, this security-related information can include any information that can be used to identify security threats. For example, the security-related information can include network-related information, such as IP addresses, domain names, asset identifiers, network traffic volume, uniform resource locator strings, and source addresses. (The process of detecting security threats for network-related information is further described in U.S. patent application Ser. Nos. 13/956,252, and 13/956,262.) Security-related information can also include endpoint information, such as malware infection data and system configuration information, as well as access control information, such as login/logout information and access failure notifications. The security-related

information can originate from various sources within a data center, such as hosts, virtual machines, storage devices and sensors. The security-related information can also originate from various sources in a network, such as routers, switches, email servers, proxy servers, gateways, firewalls and intrusion-detection systems.

During operation, the SPLUNK® APP FOR ENTERPRISE SECURITY facilitates detecting so-called “notable events” that are likely to indicate a security threat. These notable events can be detected in a number of ways: (1) an analyst can notice a correlation in the data and can manually identify a corresponding group of one or more events as “notable;” or (2) an analyst can define a “correlation search” specifying criteria for a notable event, and every time one or more events satisfy the criteria, the application can indicate that the one or more events are notable. An analyst can alternatively select a pre-defined correlation search provided by the application. Note that correlation searches can be run continuously or at regular intervals (e.g., every hour) to search for notable events. Upon detection, notable events can be stored in a dedicated “notable events index,” which can be subsequently accessed to generate various visualizations containing security-related information. Also, alerts can be generated to notify system operators when important notable events are discovered.

The SPLUNK® APP FOR ENTERPRISE SECURITY provides various visualizations to aid in discovering security threats, such as a “key indicators view” that enables a user to view security metrics of interest, such as counts of different types of notable events. For example, FIG. 77A illustrates an exemplary key indicators view 7700 that comprises a dashboard, which can display a value 7701, for various security-related metrics, such as malware infections 7702. It can also display a change in a metric value 7703, which indicates that the number of malware infections increased by 63 during the preceding interval. Key indicators view 7700 additionally displays a histogram panel 7704 that displays a histogram of notable events organized by urgency values, and a histogram of notable events organized by time intervals. This key indicators view is described in further detail in pending U.S. patent application Ser. No. 13/956,338 filed Jul. 31, 2013.

These visualizations can also include an “incident review dashboard” that enables a user to view and act on “notable events.” These notable events can include: (1) a single event of high importance, such as any activity from a known web attacker; or (2) multiple events that collectively warrant review, such as a large number of authentication failures on a host followed by a successful authentication. For example, FIG. 77B illustrates an exemplary incident review dashboard 7710 that includes a set of incident attribute fields 7711 that, for example, enables a user to specify a time range field 7712 for the displayed events. It also includes a timeline 7713 that graphically illustrates the number of incidents that occurred in one-hour time intervals over the selected time range. It additionally displays an events list 7714 that enables a user to view a list of all of the notable events that match the criteria in the incident attributes fields 7711. To facilitate identifying patterns among the notable events, each notable event can be associated with an urgency value (e.g., low, medium, high, critical), which is indicated in the incident review dashboard. The urgency value for a detected event can be determined based on the severity of the event and the priority of the system component associated with the event. The incident review dashboard is described further in “<http://docs.splunk.com/Documentation/PCI/2.1.1/User/IncidentReviewdashboard>.”

### 1.9 Data Center Monitoring

As mentioned above, the SPLUNK® ENTERPRISE platform provides various features that make it easy for developers to create various applications. One such application is the SPLUNK® APP FOR VMWARE®, which performs monitoring operations and includes analytics to facilitate diagnosing the root cause of performance problems in a data center based on large volumes of data stored by the SPLUNK® ENTERPRISE system.

This differs from conventional data-center-monitoring systems that lack the infrastructure to effectively store and analyze large volumes of performance information and log data obtained from the data center. In conventional data-center-monitoring systems, this performance data is typically pre-processed prior to being stored, for example by extracting pre-specified data items from the performance data and storing them in a database to facilitate subsequent retrieval and analysis at search time. However, the rest of the performance data is not saved and is essentially discarded during pre-processing. In contrast, the SPLUNK® APP FOR VMWARE® stores large volumes of minimally processed performance information and log data at ingestion time for later retrieval and analysis at search time when a live performance issue is being investigated.

The SPLUNK® APP FOR VMWARE® can process many types of performance-related information. In general, this performance-related information can include any type of performance-related data and log data produced by virtual machines and host computer systems in a data center. In addition to data obtained from various log files, this performance-related information can include values for performance metrics obtained through an application programming interface (API) provided as part of the vSphere Hypervisor™ system distributed by VMware, Inc. of Palo Alto, Calif. For example, these performance metrics can include: (1) CPU-related performance metrics; (2) disk-related performance metrics; (3) memory-related performance metrics; (4) network-related performance metrics; (5) energy-usage statistics; (6) data-traffic-related performance metrics; (7) overall system availability performance metrics; (8) cluster-related performance metrics; and (9) virtual machine performance statistics. For more details about such performance metrics, please see U.S. patent Ser. No. 14/167,316 filed 29 Jan. 2014, which is hereby incorporated herein by reference. Also, see “vSphere Monitoring and Performance,” Update 1, vSphere 5.5, EN-001357-00, <http://pubs.vmware.com/vsphere-55/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-551-monitoring-performance-guide.pdf>.

To facilitate retrieving information of interest from performance data and log files, the SPLUNK® APP FOR VMWARE® provides pre-specified schemas for extracting relevant values from different types of performance-related event data, and also enables a user to define such schemas.

The SPLUNK® APP FOR VMWARE® additionally provides various visualizations to facilitate detecting and diagnosing the root cause of performance problems. For example, one such visualization is a “proactive monitoring tree” that enables a user to easily view and understand relationships among various factors that affect the performance of a hierarchically structured computing system. This proactive monitoring tree enables a user to easily navigate the hierarchy by selectively expanding nodes representing various entities (e.g., virtual centers or computing clusters) to view performance information for lower-level nodes associated with lower-level entities (e.g., virtual machines or host systems). Exemplary node-expansion operations are illustrated in FIG. 77C, wherein nodes 7733 and 7734 are selectively expanded. Note that nodes 7731-7739 can be displayed using different

patterns or colors to represent different performance states, such as a critical state, a warning state, a normal state or an unknown/offline state. The ease of navigation provided by selective expansion in combination with the associated performance-state information enables a user to quickly diagnose the root cause of a performance problem. The proactive monitoring tree is described in further detail in U.S. patent application Ser. No. 14/235,490 filed on 15 Apr. 2014, which is hereby incorporated herein by reference for all possible purposes.

The SPLUNK® APP FOR VMWARE® also provides a user interface that enables a user to select a specific time range and then view heterogeneous data, comprising events, log data and associated performance metrics, for the selected time range. For example, the screen illustrated in FIG. 77D displays a listing of recent “tasks and events” and a listing of recent “log entries” for a selected time range above a performance-metric graph for “average CPU core utilization” for the selected time range. Note that a user is able to operate pull-down menus 7742 to selectively display different performance metric graphs for the selected time range. This enables the user to correlate trends in the performance-metric graph with corresponding event and log data to quickly determine the root cause of a performance problem. This user interface is described in more detail in U.S. patent application Ser. No. 14/167,316 filed on 29 Jan. 2014, which is hereby incorporated herein by reference for all possible purposes.

FIG. 78 illustrates a diagrammatic representation of a machine in the exemplary form of a computer system 7800 within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed. The system 7800 may be in the form of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed. In alternative embodiments, the machine may be connected (e.g., networked) to other machines in a LAN, an intranet, an extranet, or the Internet. The machine may operate in the capacity of a server machine in client-server network environment. The machine may be a personal computer (PC), a set-top box (STB), a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein. In one embodiment, computer system 7800 may represent system 210 of FIG. 2.

The exemplary computer system 7800 includes a processing device (processor) 7802, a main memory 7804 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) such as synchronous DRAM (SDRAM)), a static memory 7806 (e.g., flash memory, static random access memory (SRAM)), and a data storage device 7818, which communicate with each other via a bus 7830.

Processing device 7802 represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device 7802 may be a complex instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or a processor implementing other instruction sets or processors implementing a combination of instruction sets. The processing device 7802 may also be one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field program-

mable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. The processing device 7802 is configured to execute the notification manager 210 for performing the operations and steps discussed herein.

The computer system 7800 may further include a network interface device 7808. The computer system 7800 also may include a video display unit 7810 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)), an alphanumeric input device 7812 (e.g., a keyboard), a cursor control device 7814 (e.g., a mouse), and a signal generation device 7816 (e.g., a speaker).

The data storage device 7818 may include a computer-readable medium 7828 on which is stored one or more sets of instructions 7822 (e.g., instructions for search term generation) embodying any one or more of the methodologies or functions described herein. The instructions 7822 may also reside, completely or at least partially, within the main memory 7804 and/or within processing logic 7826 of the processing device 7802 during execution thereof by the computer system 7800, the main memory 7804 and the processing device 7802 also constituting computer-readable media. The instructions may further be transmitted or received over a network 7820 via the network interface device 7808.

While the computer-readable storage medium 7828 is shown in an exemplary embodiment to be a single medium, the term “computer-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term “computer-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present invention. The term “computer-readable storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, optical media, and magnetic media.

The preceding description sets forth numerous specific details such as examples of specific systems, components, methods, and so forth, in order to provide a good understanding of several embodiments of the present invention. It will be apparent to one skilled in the art, however, that at least some embodiments of the present invention may be practiced without these specific details. In other instances, well-known components or methods are not described in detail or are presented in simple block diagram format in order to avoid unnecessarily obscuring the present invention. Thus, the specific details set forth are merely exemplary. Particular implementations may vary from these exemplary details and still be contemplated to be within the scope of the present invention.

In the above description, numerous details are set forth. It will be apparent, however, to one of ordinary skill in the art having the benefit of this disclosure, that embodiments of the invention may be practiced without these specific details. In some instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the description.

Some portions of the detailed description are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities

87

take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as “determining”, “identifying”, “adding”, “selecting” or the like, refer to the actions and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (e.g., electronic) quantities within the computer system’s registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

Embodiments of the invention also relate to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions.

The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct a more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

It is to be understood that the above description is intended to be illustrative, and not restrictive. Many other embodiments will be apparent to those of skill in the art upon reading and understanding the above description. The scope of the invention should, therefore, be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

What is claimed is:

1. A method, comprising:

receiving a selection of key performance indicators (KPIs) that each indicate a different aspect of how a service provided by one or more entities is performing, wherein each entity of the one or more entities produces machine data or wherein each entity has its operation reflected in machine data not produced by the entity, wherein a stored entity definition records an association between each entity and its associated machine data, wherein a stored service definition associates the one or more entities that provide the service, wherein each KPI is defined by a different search query that derives one or more values from the machine data associated with the one or more entities providing the service, wherein each of the

88

one or more values is associated with a point in time and represents the aspect of how the service is performing at the associated point in time, and wherein the machine data comprises unstructured raw machine data; and for each of the selected KPIs:

deriving the one or more values; and

causing display of a graphical visualization of the derived one or more values for the KPI along a time-based graph lane, wherein the graph lanes for the selected KPIs are parallel to each other, and wherein the graphical visualizations in the graph lanes are all calibrated to a same time scale;

wherein the method is performed by one or more processing devices.

2. The method of claim 1, further comprising:

causing display of a timeline representing the time scale in parallel to the graph lanes.

3. The method of claim 1, wherein the graphical visualizations in the graph lanes all cover a same time range.

4. The method of claim 1, wherein the associated machine data for at least one entity comes from at least two different sources.

5. The method of claim 1, further comprising:

receiving through a graphical interface a selection of a time range for the graphical visualizations in the graph lanes.

6. The method of claim 1, wherein a time range covered by the graphical visualizations in the graph comprises a fixed period of time.

7. The method of claim 1, wherein a time range covered by the graphical visualizations in the graph comprises a rolling period of time of a set length.

8. The method of claim 1, wherein at least one of the graphical visualizations includes a line graph.

9. The method of claim 1, wherein at least one of the graphical visualizations includes an area graph.

10. The method of claim 1, wherein at least one of the graphical visualizations includes a bar chart.

11. The method of claim 1, wherein at least one of the graphical visualizations includes a heat map.

12. The method of claim 1, further comprising:

causing display of a graphical interface for selecting, from a plurality of options, a graph style for at least one of the graphical visualizations;

receiving a selection of a graph style for the at least one of the graphical visualizations; and

causing display of the at least one of the graphical visualizations in the selected graph style.

13. The method of claim 1, further comprising:

causing display, in association with a particular graph lane, a graphical interface for selecting, from a plurality of options, a graph style for displaying the graphical visualization in the associated graph lane.

14. The method of claim 1, further comprising:

receiving, through a graphical interface, a selection of a time range for the graphical visualizations in the graph lanes; and

receiving, through the graphical interface, a graph style selected from a plurality of options for displaying at least one of the graphical visualizations.

15. The method of claim 1, wherein the KPIs are selected from a list of previously created KPIs displayed in a graphical interface.

16. The method of claim 1, wherein the KPIs are selected from KPI widgets providing a numerical or graphical representation of the one or more values for corresponding KPIs displayed on a service-monitoring dashboard.

89

17. The method of claim 1, wherein at least two of the selected KPIs correspond to different services.

18. The method of claim 1, wherein at least two of the selected KPIs correspond to a same service.

19. The method of claim 1, wherein the search query is sent to an event processing system for execution.

20. The method of claim 1, further comprising:  
receiving, through a graphical interface, an indication of a title, a description and an access permission for the graphical visualizations.

21. The method of claim 1, further comprising:  
causing display of a particular search query or an identifier of the particular search query that defines at least one of the selected KPIs;

receiving input indicating that a user wants to execute the particular search query; and  
based on receiving the input, causing execution of the search query.

22. The method of claim 1, wherein the one or more values for a KPI are derived from the machine data at least in part using a late-binding schema.

23. The method of claim 1, wherein the search queries that define the selected KPIs are in a search language processed by an event processing system having a late-binding schema.

24. The method of claim 1, wherein the one or more values for a KPI are derived from time-stamped events, the time-stamped events each including at least a portion of raw machine data.

25. The method of claim 1, wherein each search query is recorded in a stored KPI definition for the KPI it defines, the KPI definition associated with the service definition for the service corresponding to that KPI.

26. The method of claim 1, wherein the search query defining at least one KPI relies upon a data model.

27. A system comprising:

a memory; and

a processing device coupled with the memory to:

receive a selection of key performance indicators (KPIs) that each indicate a different aspect of how a service provided by one or more entities is performing, wherein each entity of the one or more entities produces machine data or wherein each entity has its operation reflected in machine data not produced by the entity, wherein a stored entity definition records an association between each entity and its associated machine data, wherein a stored service definition associates the one or more entities that provide the service, wherein each KPI is defined by a different search query that derives one or more values from the machine data associated with the one or more entities providing the service, wherein each of the one or more values is associated with a point in time and represents the aspect of how the service is performing at the associated point in time, and wherein the machine data comprises unstructured raw machine data; and

90

for each of the selected KPIs:

derive the one or more values; and

cause display of a graphical visualization of the derived one or more values for the KPI along a time-based graph lane, wherein the graph lanes for the selected KPIs are parallel to each other, and wherein the graphical visualizations in the graph lanes are all calibrated to a same time scale.

28. The system of claim 27, the processing device further to:

receive, through a graphical interface, a selection of a time range for the graphical visualizations in the graph lanes; and

receive, through the graphical interface, a graph style selected from a plurality options for displaying at least one of the graphical visualizations.

29. A non-transitory computer readable storage medium encoding instructions thereon that, in response to execution by one or more processing devices, cause the processing device to perform operations comprising:

receiving a selection of key performance indicators (KPIs) that each indicate a different aspect of how a service provided by one or more entities is performing, wherein each entity of the one or more entities produces machine data or wherein each entity has its operation reflected in machine data not produced by the entity, wherein a stored entity definition records an association between each entity and its associated machine data, wherein a stored service definition associates the one or more entities that provide the service, wherein each KPI is defined by a different search query that derives one or more values from the machine data associated with the one or more entities providing the service, wherein each of the one or more values is associated with a point in time and represents the aspect of how the service is performing at the associated point in time, and wherein the machine data comprises unstructured raw machine data; and

for each of the selected KPIs:

deriving the one or more values; and

causing display of a graphical visualization of the derived one or more values for the KPI along a time-based graph lane, wherein the graph lanes for the selected KPIs are parallel to each other, and wherein the graphical visualizations in the graph lanes are all calibrated to a same time scale;

wherein the operations are performed by the one or more processing devices.

30. The non-transitory computer readable storage medium of claim 29, wherein the operations further comprise:

receiving, through a graphical interface, a selection of a time range for the graphical visualizations in the graph lanes; and

receiving, through the graphical interface, a graph style selected from a plurality options for displaying at least one of the graphical visualizations.

\* \* \* \* \*